



- Expert Verified, Online, **Free**.

HOTSPOT -

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

Hot Area:

### 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.

### Answer Area

Descriptive analytics tells you

Suggested Answer:

- 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.

Reference:

<https://demand-planning.com/2020/01/20/the-differences-between-descriptive-diagnostic-predictive-cognitive-analytics/>

- 👤 **AWSCloud** Highly Voted 3 years, 11 months ago

Descriptive analytics tells "What happened in the past". So B is the right choice.

upvoted 67 times
- 👤 **Samanouseke** Highly Voted 3 years, 5 months ago

Passed on 04/18/2021, country Brazil. 49 questions. All questions from here. Two new questions out of this dump pool. Pay attention because they changed position for the answers in some questions. A lot of drag and drop and multiple choices. Thanks Exam Topic community. Let's keep going. Best of lucky to everyone.

upvoted 46 times
- 👤 **madyjoe21** 3 years, 4 months ago

I'm from Brazil too! Nice to see u here!

upvoted 8 times
- 👤 **hismatsu3910** Most Recent 1 week, 2 days ago

The answer is B.

upvoted 1 times
- 👤 **shahrzadkhhb** 1 month ago

Descriptive analysis shows that what has happened in the past. So, the second option is correct.

upvoted 1 times
- 👤 **rapaji** 1 month, 1 week ago

B is the right choice.

upvoted 1 times
- 👤 **givaxeV** 1 month, 3 weeks ago

Achieving Microsoft DP-900 exam was straightforward with prep4certs preparation materials. They cover everything needed.

upvoted 1 times
- 👤 **AGTraining** 6 months, 2 weeks ago

option B

upvoted 1 times
- 👤 **gopolo** 10 months, 1 week ago

Our DP-900 Testing Engine will Save your DP-900 Exam Score so you can Review it later to improve your results.

<https://www.dumps4azure.com/dp-900-dumps.html>

upvoted 23 times

🗨️ 👤 **akhil5432** 1 year, 1 month ago

option B

upvoted 1 times

🗨️ 👤 **globy118** 1 year, 3 months ago

the answer is B

upvoted 1 times

🗨️ 👤 **chiiiweiii** 1 year, 6 months ago

Correct Answer

upvoted 3 times

🗨️ 👤 **Anirban\_Bose** 1 year, 9 months ago

B is the correct Answer

upvoted 4 times

🗨️ 👤 **cocolee** 2 years, 2 months ago

is this question out of syllabus of dp900 now? Seems like it should be something in dp500 now

upvoted 1 times

🗨️ 👤 **blackskies** 1 year, 11 months ago

it looks like.

upvoted 1 times

🗨️ 👤 **Gracious\_T** 2 years, 3 months ago

I just finished my exam and I passed, 929 out of 1000. Most of the questions were answered here. Thank you very much ExamTopics for making this easy for me. 🙏

upvoted 3 times

🗨️ 👤 **SamuelSami** 2 years, 5 months ago

Descriptive Analytics tells you what happened in the past. Diagnostic Analytics helps you understand why something happened in the past.

Diagnostic analytics ask about the present. They drill down into why something has happened and helps users diagnose issues. Descriptive analytics ask about the past. They want to know what has been happening to the business and how this is likely to affect future sales.

upvoted 14 times

🗨️ 👤 **eusebio72** 2 years, 5 months ago

Passed thisd morning. 49 questions and 46 present here. 3 are new

upvoted 3 times

🗨️ 👤 **AZ5cert** 2 years, 5 months ago

Correct Answer: B

A is "Predictive"

C is "Prescriptive"

D is "Diagnostic"

upvoted 14 times

HOTSPOT -

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE: Each correct selection is worth one point.

Hot Area:

### Answer Area

Statements	Yes	No
Normalization involves eliminating relationships between database tables.	<input type="radio"/>	<input type="radio"/>
Normalizing a database reduces data redundancy.	<input type="radio"/>	<input type="radio"/>
Normalization improves data integrity.	<input type="radio"/>	<input type="radio"/>

### Answer Area

Statements	Yes	No
Normalization involves eliminating relationships between database tables.	<input type="radio"/>	<input checked="" type="radio"/>
Normalizing a database reduces data redundancy.	<input checked="" type="radio"/>	<input type="radio"/>
Normalization improves data integrity.	<input checked="" type="radio"/>	<input type="radio"/>

Reference:

<https://www.sqlshack.com/what-is-database-normalization-in-sql-server/>

 **giventheman** Highly Voted 3 years, 11 months ago

no yes yes

upvoted 61 times

 **yeanlingmedal71** Highly Voted 3 years, 4 months ago

Normalization is the process of organizing data in a database. This includes creating tables and establishing relationships between those tables according to rules designed both to protect the data and to make the database more flexible by eliminating redundancy and inconsistent dependency.

[https://docs.microsoft.com/en-us/office/troubleshoot/access/database-normalization-](https://docs.microsoft.com/en-us/office/troubleshoot/access/database-normalization-description#:~:text=Normalization%20is%20the%20process%20of,eliminating%20redundancy%20and%20inconsistent%20dependency)

[description#:~:text=Normalization%20is%20the%20process%20of,eliminating%20redundancy%20and%20inconsistent%20dependency](https://docs.microsoft.com/en-us/office/troubleshoot/access/database-normalization-description#:~:text=Normalization%20is%20the%20process%20of,eliminating%20redundancy%20and%20inconsistent%20dependency)

upvoted 19 times

 **hismatsu3910** Most Recent 1 week, 2 days ago

NO-YES-YES

upvoted 1 times

 **shahrzadkhh** 1 month ago

the aim of the database normalization is to reduce the data redundancy and data integrity. So, the answer would be No, Yes, Yes.

upvoted 1 times

 **mnithi** 3 months, 3 weeks ago

I saw this question in yesterdays exam

upvoted 1 times

 **AGTraining** 6 months, 2 weeks ago

No,Yes,Yes

upvoted 1 times

 **trashbox** 10 months, 2 weeks ago

No, Yes, Yes.

Appeared on 2023-11-05.

upvoted 1 times

  **akhil5432** 1 year, 1 month ago

No

Yes

Yes

upvoted 1 times

  **VNR25** 1 year, 2 months ago

This question appeared in the 07/07/23 exam and the answer is correct.

upvoted 2 times

  **Vanessa23** 1 year, 5 months ago

Normalization is tables related PK-FK, reduces space. Normalization is recommended for relational databases.

upvoted 1 times

  **chiiiweiii** 1 year, 6 months ago

YNN are the Correct Answer

upvoted 1 times

  **chiiiweiii** 1 year, 6 months ago

NY...sorry

upvoted 3 times

  **Joanna0** 1 year, 6 months ago

Correct!

NY

upvoted 1 times

  **King\_Lam** 1 year, 7 months ago

In exam 27/1/23

upvoted 2 times

  **axgabi** 1 year, 9 months ago

in Exam 18 dec 2022

upvoted 2 times

  **Ozzy111** 2 years, 4 months ago

N, Y, Y

upvoted 1 times

  **heenak55** 2 years, 4 months ago

In exam 21 April 2022

upvoted 2 times

  **AZ5cert** 2 years, 5 months ago

Correct Answer: N Y Y

upvoted 2 times

HOTSPOT -

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

Hot Area:

### Answer Area

An extract, transform, and load (ETL) process requires

- a matching schema in the data source and the data target.
- a target data store powerful enough to transform data.
- data that is fully processed before being loaded to the target data store.
- that the data target be a relational database.

Suggested Answer:

### Answer Area

An extract, transform, and load (ETL) process requires

- a matching schema in the data source and the data target.
- a target data store powerful enough to transform data.
- data that is fully processed before being loaded to the target data store.
- that the data target be a relational database.

Extract, transform, and load (ETL) is a data pipeline used to collect data from various sources, transform the data according to business rules, and load it into a destination data store.

Reference:

<https://docs.microsoft.com/en-us/azure/architecture/data-guide/relational-data/etl>

 **muks1003** Highly Voted 2 years, 2 months ago

For ETL it should be C, for ELT it will be B.

upvoted 28 times

 **heenak55** Highly Voted 2 years, 4 months ago

I am confused for the answer why it is C instead of B

upvoted 5 times

 **pads98** 2 years, 4 months ago

ETL process, note the order EXTRACT, TRANSFORM, LOAD.

data should be TRANSFORMED(processed) before it is loaded.

upvoted 14 times

 **dongboy** 2 years, 4 months ago

It's in the name, E T L. The data needs to be transformed before it is loaded. If this was ELT, then the answer would be B.

upvoted 5 times

 **shahrzadkhb** Most Recent 1 month ago

ETL stands for extract, transform, and load. So, data must be transformed before loading. Therefore, the third option is correct.

upvoted 1 times

 **Jstein91** 1 month, 1 week ago

Answer is correct -C. - BEFORE the Data store - Extract, load, transform (ELT) differs from ETL solely in where the transformation takes place. In the ELT pipeline, the transformation occurs in the target data store. - This question asks - before being loaded not "in" the data store - Which is ETL

upvoted 1 times

 **Abhishek17** 3 months, 4 weeks ago

No yes yes

upvoted 1 times

🗨️ 👤 **AGTraining** 6 months, 2 weeks ago

option C

upvoted 2 times

🗨️ 👤 **MannyJZ** 1 year, 1 month ago

On exam 08/08/23

upvoted 2 times

🗨️ 👤 **akhil5432** 1 year, 1 month ago

opt C is correct answer

upvoted 1 times

🗨️ 👤 **HRKK** 1 year, 2 months ago

i think it is the correct answer

upvoted 1 times

🗨️ 👤 **fifulka** 1 year, 3 months ago

Was on the exam 16/06/2023

upvoted 2 times

🗨️ 👤 **Vanessa23** 1 year, 5 months ago

ETL - extract, transform (before load) and load. target datastore only load data (no transformation)

upvoted 1 times

🗨️ 👤 **Dudures** 1 year, 5 months ago

Exam 04/15/2023

upvoted 2 times

🗨️ 👤 **chiiweiii** 1 year, 6 months ago

Correct Answer

ELT-a target data store powerful enough to transform data.

ETL-data that is fully processed before being loaded to the target data store.

upvoted 3 times

🗨️ 👤 **Rakrah** 1 year, 8 months ago

ETL - Extract then Transform the Load - It always ensure apply all the transformation logic after extracting/bore loading into Target table. Target table is ready to use the data for reporting/visualizations. My preference is "C"

upvoted 1 times

🗨️ 👤 **KMLearn3** 1 year, 9 months ago

Pretty similar to Qst#31 (ELT there). Here ETL.

It will be understandable if you take "transform" and "process" as equivalent.

And then yes, process/transform is before load.

And on ELT (Qst#31): yes, for transform/process the target needs capabilities for it.

upvoted 1 times

🗨️ 👤 **KMLearn3** 1 year, 9 months ago

Missleading because question is about "requires" (=prerequisite) and then the answer "fully processed before". Extraction belongs to processing (ETL or ELT). So why should that have startet after "FULLY processed before"?

upvoted 1 times

🗨️ 👤 **edurakhan** 1 year, 10 months ago

Exam 11/15/22

upvoted 1 times

HOTSPOT -

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

Hot Area:

## Answer Area

In batch processing,

- data is always inserted one row at a time.
- data is processed in real-time.
- latency in delivering data processing results is acceptable.
- processing can only execute serially.

## Answer Area

Suggested Answer: In batch processing,

- data is always inserted one row at a time.
- data is processed in real-time.
- latency in delivering data processing results is acceptable.
- processing can only execute serially.

When to use batch processing.

You might expect latencies when using batch processing. For many situations, however, this type of delay before the transfer of data begins is not a big issue - the processes that use this function are not mission critical at the exact moment.

Reference:

<https://www.bmc.com/blogs/what-is-batch-processing-batch-processing-explained/>

 **fusioner** Highly Voted 2 years ago

Correct Answer. Latency is expected in Batch Processing  
upvoted 10 times

 **dataalex** Highly Voted 8 months ago

Hi everyone! I just passed exam, my score is more then 900.

1. Exam 80-90% consisted from questions from this site.
  2. Answers on this site for this questions are not always correct, and comments with answers are not always correct.
  3. If you are reading in comments "I passed exam and this question was on exam, the answers is \_\_ (gives you some answer)" - dont trust it blindly! They just tells their opinion about the answer. Because when you recieve you certification, your score, you do not recieve informations on what questions exactly you answered right or wrong! You just know your score and your score on every section. You dont know the answer and you no longer have access to question that you answered.
  - 4) Everyone good luck! You can do it.
- upvoted 7 times

 **shahrzadkhb** Most Recent 1 month ago

I believe the given answer is correct.

Resource:

<https://learn.microsoft.com/en-us/training/modules/explore-fundamentals-stream-processing/2-batch-stream>  
upvoted 1 times

 **rapaji** 1 month, 1 week ago

Option C

upvoted 1 times

 **AGTraining** 6 months, 2 weeks ago

option C

upvoted 1 times

🗨️ 👤 **BeliteLiu** 1 year, 2 months ago

I acceding for exam on 2023/7/11, hope i can get good score.

upvoted 2 times

🗨️ 👤 **Vanessa23** 1 year, 5 months ago

Batch = latency ~ minutes or hours

upvoted 4 times

🗨️ 👤 **ahmedkmj** 1 year, 5 months ago

in the exam 26/03

upvoted 2 times

🗨️ 👤 **chiiiweiii** 1 year, 6 months ago

Correct Answer.

upvoted 2 times

🗨️ 👤 **Abhi256** 1 year, 9 months ago

Latency is expected.

upvoted 4 times

🗨️ 👤 **captnfabinou** 1 year, 9 months ago

was on exam on 12/08/2022

upvoted 3 times

🗨️ 👤 **blackskies** 1 year, 11 months ago

was on exam today 09/30/22

upvoted 4 times

🗨️ 👤 **zekblue** 2 years ago

correct answer

upvoted 4 times

HOTSPOT -

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

Hot Area:

### Answer Area

Transcribing audio files is an example of  analytics.

<input type="text"/>
cognitive
descriptive
predictive
prescriptive

### Answer Area

Suggested Answer:

Transcribing audio files is an example of  analytics.

<input type="text"/>
cognitive
descriptive
predictive
prescriptive

Reference:

<https://azure.microsoft.com/en-us/services/cognitive-services/speech-services/>

**jasu** Highly Voted 4 years ago

Transcribe means converting from Audio to Text. This is similar to Alexa, Alexa is a Cognitive device. Cognitive is the correct answer.  
upvoted 144 times

**Tremayne** 2 years, 1 month ago

Great explanation  
upvoted 5 times

**Sc2** Highly Voted 4 years ago

I agree with cognitive  
upvoted 40 times

**AGTraining** Most Recent 6 months, 2 weeks ago

option A  
upvoted 1 times

**Mega2** 7 months ago

Transcribing audio files is an example of cognitive processing.

Cognitive processing involves understanding, interpreting, and deriving meaning from data. In the case of transcribing audio files, cognitive processing algorithms are used to analyze and interpret the spoken words in the audio and convert them into text. This process requires understanding the context, language, and nuances of human speech, which falls under the category of cognitive tasks.

upvoted 3 times

**akhil5432** 1 year, 1 month ago

Cognitive  
upvoted 1 times

**srgopalkr** 1 year, 1 month ago

Answer is Cognitive, given the speech-to-text service in Azure is classified as under Cognitive Services.  
upvoted 1 times

**AKATTHULA** 1 year, 4 months ago

This is very tricky, everywhere this answer is marked for Cognitive but Chatgpt says it's descriptive because the nature of analytics is based on analyzing from the past, so it's descriptive and not cognitive.  
upvoted 2 times

AKATTHULA 1 year, 4 months ago

Correction: as of 5/1 chatgpt shows as cognitive as well. Cognitive is correct.

upvoted 1 times

AKATTHULA 1 year, 4 months ago

There is another question similar to this that looks like: creating closed caption text for audio files = which analytics. Answer is Descriptive in this case.

upvoted 1 times

navbizmd 1 year, 5 months ago

Transcribing audio files can be considered as a form of descriptive analytics. Descriptive analytics is a type of analytics that focuses on summarizing historical data to gain insights into past events and trends. In the case of transcribing audio files, the process involves converting spoken language into written text, which can then be analyzed to understand what was said in the past.

Cognitive analytics, on the other hand, is a type of analytics that involves using machine learning algorithms to analyze unstructured data such as text, images, and speech. Cognitive analytics is focused on identifying patterns and insights in data that are not immediately apparent to humans, and it is often used for predictive and prescriptive analytics.

upvoted 3 times

chiiiweiii 1 year, 6 months ago

Correct Answer.

upvoted 1 times

Alex4190 1 year, 6 months ago

Only showing 5 questions per page to force non contributors to do more annoying captcha? Great way to push someone to find and use a different site.....

Cognitive is correct

upvoted 4 times

PatYeo 1 year, 11 months ago

<https://demand-planning.com/2020/01/20/the-differences-between-descriptive-diagnostic-predictive-cognitive-analytics/>

Cognitive is correct.

upvoted 4 times

GURU112021 2 years, 2 months ago

Answer is correct

upvoted 1 times

abosafi87 2 years, 4 months ago

cognitive is the correct answer

upvoted 1 times

Atemius 2 years, 5 months ago

Cognitive

<https://docs.microsoft.com/en-us/azure/cognitive-services/speech-service/batch-transcription>

upvoted 3 times

AZ5cert 2 years, 5 months ago

Correct Answer: Cognitive

upvoted 1 times

RajeshAzure 2 years, 6 months ago

correct answer is C

upvoted 1 times

yomar 2 years, 7 months ago

LA respuesta es la A-cognitive

upvoted 1 times

DRAG DROP -

Match the types of analytics that can be used to answer the business questions.

To answer, drag the appropriate analytics type from the column on the left to its question on the right. Each analytics type may be used once, more than once, or not at all.

NOTE: Each correct match is worth one point.

Select and Place:

### Analytics Types

Cognitive

Diagnostic

Descriptive

Predictive

Prescriptive

### Answer Area

Why did sales increase last month?

How do I allocate my budget to buy different inventory items?

Which people are mentioned in a company's business documents?

#### Suggested Answer: *Incorrect Answer:*

Box 1: Diagnostic -

Diagnostic Analytics: At this stage you can begin to answer some of those why questions. Historical data can begin to be measured against other data to answer the question of why something happened in the past. This is the process of gathering and interpreting different data sets to identify anomalies, detect patterns, and determine relationships.

Box 2: Prescriptive -

Prescriptive analytics is a combination of data, mathematical models, and various business rules to infer actions to influence future desired outcomes.

Predictive analytics, broadly speaking, is a category of business intelligence that uses descriptive and predictive variables from the past to analyze and identify the likelihood of an unknown future outcome

Box 3: Descriptive -

⇒ 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?



Reference:

<https://demand-planning.com/2020/01/20/the-differences-between-descriptive-diagnostic-predictive-cognitive-analytics/>

<https://azure.microsoft.com/en-us/blog/answering-whats-happening-whys-happening-and-what-will-happen-with-iot-analytics/>

WillyMac Highly Voted 2 years, 11 months ago

The last must be "Cognitive".

That's what about Azure Cognitive Services

upvoted 45 times

floop92 2 years, 8 months ago

Agreed! Had this in the course

upvoted 8 times

gs23mi Highly Voted 2 years, 11 months ago

1 diagnostic : it answers the question "why did it happen?"

2 prescriptive : it answers the question "how can we make it happen?"

3 descriptive: it answers the question "what happened"

ref: Microsoft Azure Data Fundamentals - page 23-24

upvoted 43 times

Yaizapez 2 years, 7 months ago

The third scenario doesn't look into "what happened": It's asking for text analytics, which is not a descriptive function but a cognitive one.

upvoted 35 times

AGTraining Most Recent 6 months, 2 weeks ago

Diagnostic

Prescriptive

Cognitive

upvoted 1 times

BXtheCoder 7 months, 2 weeks ago

Isn't box 3 cognitive?

upvoted 1 times

akhil5432 1 year, 1 month ago

diagnostic

prescriptive

discriptive

upvoted 1 times

fiksarion 1 year, 2 months ago

Justification:

Diagnostic analytics is used to understand the cause and effect relationships in data, which makes it suitable for answering the question about why sales increased last month. It helps identify the factors or variables that contributed to the increase in sales.

Prescriptive analytics focuses on providing recommendations or optimal solutions to business problems. Allocating a budget to buy different inventory items requires making decisions based on various factors and constraints, making prescriptive analytics the appropriate type.

Cognitive analytics involves analyzing unstructured data such as text, images, or audio to gain insights. Identifying people mentioned in business documents would require analyzing text data, making cognitive analytics the relevant type in this case.

upvoted 4 times

Rajivjain 1 year, 3 months ago

Correct Answer:

The statement "Which people are mentioned in a company's business documents?" is descriptive. It is asking for a description or explanation of the individuals that are typically mentioned in a company's business documents. It does not involve any cognitive or analytical processing.

upvoted 2 times

okayhey 6 months, 1 week ago

That's not what descriptive analytics is, descriptive analysis looks at the past

upvoted 1 times

Meghwanshi 1 year, 2 months ago

Did you get this correct in the exam?

upvoted 1 times

Vanessa23 1 year, 5 months ago

Diagnostic - Prescriptive - Cognitive

upvoted 4 times

XtraWest 1 year, 5 months ago

correct answer

upvoted 2 times

chiiweiii 1 year, 6 months ago

Correct Answer.

upvoted 1 times

🗨️ **amit3770** 1 year, 8 months ago

3 rd one seems cognitive to me

upvoted 4 times

🗨️ **KMLearn3** 1 year, 9 months ago

Page not found:

<https://demand-planning.com/2020/01/20/the-differences-between-descriptive-diagnostic-predictive-cognitive-analytics/%20https://azure.microsoft.com/en-us/blog/answering-whats-happening-whys-happening-and-what-will-happen-with-iot-analytics/>

upvoted 1 times

🗨️ **rohjay** 1 year, 11 months ago

The third one is incorrect, it should be cognitive as it searches the company employees' names from some documents.

upvoted 8 times

🗨️ **JA2018** 1 year, 11 months ago

Agreed. You can refer to this.

What is Named Entity Recognition (NER) in Azure Cognitive Service for Language?

<https://learn.microsoft.com/en-us/azure/cognitive-services/language-service/named-entity-recognition/overview>

upvoted 4 times

🗨️ **RajeshAzure** 2 years, 6 months ago

correct answer

upvoted 2 times

🗨️ **at\_me\_jay** 2 years, 5 months ago

Could you please explain why? It's not really a description of anything right?

upvoted 1 times

🗨️ **yomar** 2 years, 7 months ago

-diagnostica

-Prescriptive

-cognitivo

upvoted 8 times

🗨️ **JustinTom** 2 years, 7 months ago

Diagnostic

Prescriptive

Cognitive

upvoted 6 times

🗨️ **rick245** 2 years, 7 months ago

Shouldn't 3) be cognitive?

<https://docs.microsoft.com/en-us/learn/modules/analyze-text-with-text-analytics-service/1-introduction>

In Microsoft Azure, the Language cognitive service can help simplify application development by using pre-trained models that can:

...

Identify and categorize entities in the text. Entities can be people, places, organizations, or even everyday items such as dates, times, quantities, and so on.

The question asks "Which people are mentioned in a company's business documents?" and that is exactly an application of cognitive services

In this module, you'll explore some of these capabilities and gain an understanding of how you might apply them to applications such as:

...

A tool to extract brand information or company names from documents or other text for identification purposes.

upvoted 8 times

🗨️ **BabyBee** 2 years, 5 months ago

Indeed 3 is Cognitive.

Identify and categorize entities in the text. Entities can be people, places, organizations, or even everyday items such as dates, times,

quantities, and so on.

upvoted 2 times

HOTSPOT -

You have the following JSON document.

```
{
  "customer" : {
    "first name" : "Ben",
    "last name" : "Smith",
    "address" : {
      "line 1" : "161 Azure Ln",
      "line 2" : "Palo Alto",
      "ZIP code" : "54762"
    },
    "social media": [
      {
        "service" : "twitter",
        "handle" : "@bensmith"
      },
      {
        "service" : "linkedin",
        "handle" : "bensmith"
      }
    ],
    "phone numbers": [
      {
        "type" : "mobile",
        "number" : "555-555-555"
      }
    ]
  }
}
```

Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the JSON document.

NOTE: Each correct selection is worth one point.

Hot Area:

### Answer Area

Customer is [answer choice].

	▼
a nested array	
a nested object	
a root object	

Address is [answer choice].

	▼
a nested array	
a nested object	
a root object	

Social media is [answer choice].

	▼
a nested array	
a nested object	
a root object	

## Answer Area

Customer is [answer choice].

	▼
a nested array	
a nested object	
a root object	

Address is [answer choice].

	▼
a nested array	
a nested object	
a root object	

Suggested Answer:

Social media is [answer choice].

	▼
a nested array	
a nested object	
a root object	

Reference:

[https://www.w3schools.com/js/js\\_json\\_arrays.asp](https://www.w3schools.com/js/js_json_arrays.asp)

[https://www.w3schools.com/js/js\\_json\\_objects.asp](https://www.w3schools.com/js/js_json_objects.asp)

PM2 Highly Voted 3 years, 11 months ago

Correct Answer

upvoted 49 times

Armandoo Highly Voted 3 years, 3 months ago

Root Object, Nested Root Object, and Array because of []

upvoted 22 times

AGTraining Most Recent 6 months, 2 weeks ago

root object

nested object

nested array

upvoted 2 times

MannyJZ 1 year, 1 month ago

On exam 08/08/23

upvoted 1 times

akhil5432 1 year, 1 month ago

root abject

nested object

nested array

upvoted 2 times

PsiKx 1 year, 2 months ago

Correct it is

upvoted 1 times

chiiweiii 1 year, 6 months ago

Correct Answer.

upvoted 2 times

Eltooth 2 years, 2 months ago

Root

Object

Array

upvoted 3 times

D\_Mo\_ZY 2 years, 6 months ago

Correct Answer

Was in the exam (18/03/2022)

upvoted 9 times

🗉 👤 **RajeshAzure** 2 years, 6 months ago

right answer

upvoted 1 times

🗉 👤 **yomar** 2 years, 7 months ago

LA respuesta es la C-B-A

upvoted 1 times

🗉 👤 **sumanshu** 2 years, 7 months ago

Root, Nested Object, Nested Array

upvoted 2 times

🗉 👤 **Aneta\_ane** 2 years, 9 months ago

appeared on the exam on 27/11/2021

upvoted 5 times

🗉 👤 **theLazyGuy** 2 years, 10 months ago

SocialMedia is an Array object, as someone underline due to [ ], references:

<https://docs.microsoft.com/en-us/answers/questions/200268/adf-reading-array-of-objects-in-data-flow.html>

<https://docs.microsoft.com/en-us/azure/data-factory/format-json>

upvoted 5 times

🗉 👤 **dumbmf** 3 years, 1 month ago

sahi hai bhai tension mat lo ok

upvoted 3 times

🗉 👤 **Alexpet** 3 years, 3 months ago

Social media is the same as address in structure therefore it is a Nested Object. Arrays are usually sequences of numbers whereas elements in Address and Social Media are key and value types.

upvoted 2 times

🗉 👤 **AVP\_Riga** 3 years, 3 months ago

JSON array is between [].

upvoted 4 times

🗉 👤 **mattvasc** 3 years, 3 months ago

Social Media is an array of objects. You can check the syntax in json.org website.

upvoted 4 times

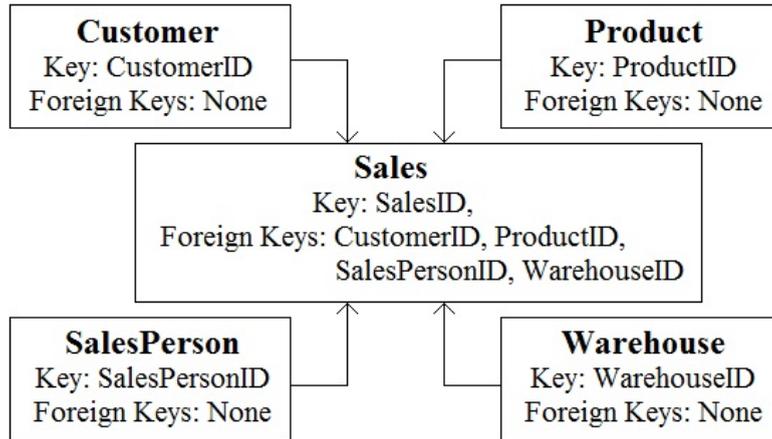
🗉 👤 **ranimswami** 3 years, 3 months ago

customer is root obj

upvoted 1 times

HOTSPOT -

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

Hot Area:

**Answer Area**

The data model is a [answer choice].

▼
transactional model
star schema
snowflake schema

Customer is a [answer choice] table.

▼
fact
dimension
bridge

**Answer Area**

The data model is a [answer choice].

▼
transactional model
star schema
snowflake schema

Suggested Answer:

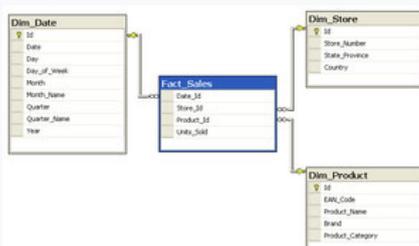
Customer is a [answer choice] table.

▼
fact
dimension
bridge

Box 1: star schema -

In computing, the star schema is the simplest style of data mart schema and is the approach most widely used to develop data warehouses and dimensional data marts. The star schema consists of one or more fact tables referencing any number of dimension tables. The star schema is an important special case of the snowflake schema, and is more effective for handling simpler queries.

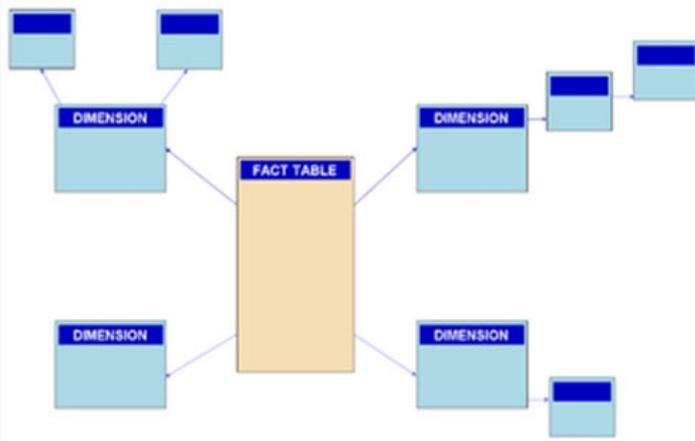
Example:



Incorrect Answers:

The data in the question is not normalized.

The snowflake schema is a variation of the star schema, featuring normalization of dimension tables. Example:



Note: A snowflake schema is a logical arrangement of tables in a multidimensional database such that the entity relationship diagram resembles a snowflake shape. The snowflake schema is represented by centralized fact tables which are connected to multiple dimensions. [citation needed]. "Snowflaking" is a method of normalizing the dimension tables in a star schema. When it is completely normalized along all the dimension tables, the resultant structure resembles a snowflake with the fact table in the middle.

Box 2: dimension -

The star schema consists of one or more fact tables referencing any number of dimension tables.

Reference:

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-tables-overview>

[https://en.wikipedia.org/wiki/Star\\_schema](https://en.wikipedia.org/wiki/Star_schema) [https://en.wikipedia.org/wiki/Snowflake\\_schema](https://en.wikipedia.org/wiki/Snowflake_schema) <https://azure.microsoft.com/en-us/blog/data-models-within-azure-analysis-services-and-power-bi/>

**chaudha4** Highly Voted 3 years, 5 months ago

Seems like this question is outside the scope of DP-900 exam. It is not covered in Azure training for DP-900 exam.  
upvoted 26 times

**tdasuni001** 3 years, 5 months ago

Agreed  
upvoted 3 times

**RajeshAzure** 2 years, 6 months ago

I saw it in MS Tutorial video on DP-900, so i assume it is valid question.  
upvoted 2 times

**pads98** 2 years, 4 months ago

This was covered in the Microsoft dp900 learning path.  
Answer is correct.  
upvoted 6 times

**DrCloud** 3 years, 2 months ago

Correct Answer.

<https://docs.microsoft.com/en-us/azure/architecture/guide/technology-choices/data-store-overview#column-family-databases>

Under "Data analytics" section:

Data type

Usually denormalized in a "star" or "snowflake" schema, consisting of fact and dimension tables.

upvoted 6 times

**yeanningmedal71** Highly Voted 3 years, 3 months ago

Star Schema in data warehouse, in which the center of the star can have one fact table and a number of associated dimension tables. It is known as star schema as its structure resembles a star.

Dimension tables describe business entities—the things you model. Entities can include products, people, places, and concepts including time itself.

Fact tables store observations or events, and can be sales orders, stock balances, exchange rates, temperatures, etc. A fact table contains dimension key columns that relate to dimension tables, and numeric measure columns.

<https://docs.microsoft.com/en-us/power-bi/guidance/star-schema>

upvoted 23 times

🗨️ 👤 **AGTraining** Most Recent 6 months, 2 weeks ago

Star

Dimension

upvoted 2 times

🗨️ 👤 **chiiiweiii** 1 year, 6 months ago

Correct Answer.

upvoted 2 times

🗨️ 👤 **xRiot007** 2 years, 2 months ago

Star schema - you have multiple dimension tables linked by a fact table.

Customer is one of the dimension tables.

upvoted 7 times

🗨️ 👤 **Eltooth** 2 years, 2 months ago

Star

Dimension

upvoted 1 times

🗨️ 👤 **Kermedi2018** 2 years, 4 months ago

It was on 17/03/2022

upvoted 3 times

🗨️ 👤 **D\_Mo\_ZY** 2 years, 6 months ago

Correct Answer

Was in the exam (18/03/2022)

upvoted 5 times

🗨️ 👤 **cj\_kuo** 2 years, 6 months ago

Answer was correct due to <https://docs.microsoft.com/en-us/power-bi/guidance/star-schema>

upvoted 3 times

🗨️ 👤 **yomar** 2 years, 7 months ago

la respuesta es

-modelo copo de nieve

-dimension

upvoted 1 times

🗨️ 👤 **sumanshu** 2 years, 7 months ago

Star Schema

Dimension

upvoted 1 times

🗨️ 👤 **Charly0710** 2 years, 8 months ago

La respuesta es STAR, No Es Snowflake.

<https://todobi.com/star-y-snowflake-schemas-que/>

upvoted 2 times

🗨️ 👤 **Aneta\_ane** 2 years, 9 months ago

appeared on the exam on 27/11/2021

upvoted 6 times

🗨️ 👤 **CyLa** 3 years ago

Fact --> Dimensions Then Dimensions become Facts and break into more Dimensions to form Snowflakes, then more of these snowflakes to become Galaxy.

upvoted 2 times

🗨️ 👤 **Armandoo** 3 years, 3 months ago

Model: Star Schema

Customer: Dimension

upvoted 6 times

🗨️ 👤 **mar765** 3 years, 3 months ago

What is it that tells us that this is not a Transactional Model?

upvoted 1 times

🗨️ 👤 **BertieT** 3 years, 1 month ago

An application may further normalize a sales transaction into header and lines. There may also be a status. This lacks a status -- so is merely expressing the sales of a Product to a Customer by Sales Person at a Warehouse -- that is, more like a statement of a "fact" after it has already occurred.

upvoted 1 times

🗨️ 👤 **DeepMoon** 3 years, 4 months ago

Customer is the Root Object.

Address is an object inside the Customer; so it a nested object.

Social media is an array with two elements inside of it (twitter/linkedin info).

So the answer is correct.

upvoted 3 times

HOTSPOT -

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

Hot Area:

### 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. |   |

### Answer Area

The massively parallel processing (MPP) engine of Azure

Synapse Analytics

Suggested Answer:

- |  |   |
|--|---|
|  | ▼ |
| distributes processing across compute nodes.       |   |
| distributes processing across control nodes.       |   |
| redirects client connections across compute nodes. |   |
| redirects client connections across control nodes. |   |

Reference:

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/massively-parallel-processing-mpp-architecture>

🗋️ **hmAZtime** Highly Voted 3 years, 5 months ago

It's A.

upvoted 51 times

🗋️ **valente\_sven1** 3 years ago

agree. Yes it should be A.

upvoted 6 times

🗋️ **sreekarv** Highly Voted 3 years, 9 months ago

The correct answer is "across compute nodes". Distributions map to Compute nodes for processing. As you pay for more compute resources, distributions are remapped to available Compute nodes.

Reference: <https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/massively-parallel-processing-mpp-architecture> [see the architecture diagram]

upvoted 14 times

🗋️ **mnithi** Most Recent 3 months, 3 weeks ago

this question appeared in my exam yesterday

upvoted 1 times

🗋️ **AGTraining** 6 months, 2 weeks ago

option A

upvoted 1 times

🗋️ **DavidPham1997** 6 months, 4 weeks ago

MPP systems are designed to be distributed parallel processing solutions, meaning they are not only able to scale up by adding more compute resources but can also scale out by adding more nodes to the system.

upvoted 1 times

🗋️ **trashbox** 10 months, 2 weeks ago

A.

Appeared on 2023-11-05.

upvoted 2 times

🗨️ 👤 **Nino\_Sasania** 11 months, 1 week ago

This question appeared in the 12/10/23 exam and the answer is correct.  
upvoted 1 times

🗨️ 👤 **VNR25** 1 year, 2 months ago

This question appeared in the 07/07/23 exam and the answer is correct.  
upvoted 2 times

🗨️ 👤 **globy118** 1 year, 2 months ago

Was on exam 24/6/2023  
upvoted 3 times

🗨️ 👤 **Dudures** 1 year, 5 months ago

Exam 04/15/2023  
upvoted 3 times

🗨️ 👤 **mamahani** 1 year, 5 months ago

I took the exam 12.04.2023 and got 1000, 100% score; this question appeared, correct answer:  
distributes processing across compute nodes  
upvoted 1 times

🗨️ 👤 **ahmedkmj** 1 year, 5 months ago

in the exam 26/03  
upvoted 1 times

🗨️ 👤 **vizviz** 1 year, 6 months ago

Was on exam March 9, 2023  
upvoted 1 times

🗨️ 👤 **chiiiweiii** 1 year, 6 months ago

Correct Answer.  
upvoted 1 times

🗨️ 👤 **King\_Lam** 1 year, 7 months ago

Correct. Was on exam 27/1/2023  
upvoted 1 times

🗨️ 👤 **captmfabinou** 1 year, 9 months ago

was on exam on 12/08/2022  
upvoted 1 times

🗨️ 👤 **blackskies** 1 year, 11 months ago

was on exam today 09/30/22  
upvoted 2 times

HOTSPOT -

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

Hot Area:

### Answer Area

	▼
A clustered index	
A FileTable	
A foreign key	
A stored procedure	

is an object associated with a table that sorts and stores the data rows in the table based on their key values.

### Answer Area

Suggested Answer:

	▼
A clustered index	
A FileTable	
A foreign key	
A stored procedure	

is an object associated with a table that sorts and stores the data rows in the table based on their key values.

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/indexes/clustered-and-nonclustered-indexes-described?view=sql-server-ver15>

🗨️ **Chandra111** Highly Voted 3 years, 6 months ago

Cluster Index.

upvoted 44 times

🗨️ **yeanlingmedal71** Highly Voted 3 years, 3 months ago

Clustered indexes sort and store the data rows in the table or view based on their key values. These are the columns included in the index definition. There can be only one clustered index per table, because the data rows themselves can be stored in only one order.

upvoted 32 times

🗨️ **AGTraining** Most Recent 6 months, 2 weeks ago

Option A

upvoted 1 times

🗨️ **Anuragsharma83** 11 months, 1 week ago

Correct Answer

upvoted 1 times

🗨️ **MannyJZ** 1 year, 1 month ago

On exam 08/08/23

upvoted 1 times

🗨️ **globy118** 1 year, 2 months ago

Was on exam 24/6/2023

upvoted 2 times

🗨️ **PhyoPluto** 1 year, 6 months ago

Appeared on the exam on 07/03/2023

upvoted 4 times

🗨️ **chiiweiii** 1 year, 6 months ago

Correct Answer.

upvoted 1 times

🗨️ **ThomasDB** 1 year, 11 months ago

Exam: 12/10/2022

upvoted 4 times

🗨️ **ria97** 2 years, 1 month ago

Was in exam 11/08. Correct answer  
upvoted 4 times

🗨️ 👤 **Drchattss** 2 years, 1 month ago  
01/08.....EXAM  
upvoted 4 times

🗨️ 👤 **Breeze21** 2 years, 2 months ago  
Exam: 19.07.22  
upvoted 3 times

🗨️ 👤 **Eltooth** 2 years, 2 months ago  
Clustered index  
upvoted 2 times

🗨️ 👤 **dk87** 2 years, 4 months ago  
Exam: 05/18/2022  
upvoted 2 times

🗨️ 👤 **purvalstrikes** 2 years, 4 months ago  
Exam: 05/12/2022  
upvoted 2 times

🗨️ 👤 **AZ5cert** 2 years, 5 months ago  
Correct Answer: A, Clustered Index  
upvoted 2 times

🗨️ 👤 **Alpyas** 2 years, 5 months ago  
Clustered indexes sort and store the data rows in the table or view based on their key values. These are the columns included in the index definition. There can be only one clustered index per table, because the data rows themselves can be stored in only one order. Nonclustered indexes have a structure separate from the data rows. A nonclustered index contains the nonclustered index key values and each key value entry has a pointer to the data row that contains the key value.  
upvoted 5 times

HOTSPOT -

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

Hot Area:

### Answer Area

A relational database is appropriate for scenarios that involve a high volume of

	▼
changes to relationships between entities	
geographically distributed writes	
transactional writes	
writes that have varying data structures	

Suggested Answer:

### Answer Area

A relational database is appropriate for scenarios that involve a high volume of

	▼
changes to relationships between entities	
geographically distributed writes	
transactional writes	
writes that have varying data structures	

Disadvantages of non-relational databases include: Data Consistency & " non-relational databases do not perform ACID transactions.

Note: Relational databases are optimized for writes. They are optimized for consistency and availability. Advantages of relational databases include simplicity, ease of data retrieval, data integrity, and flexibility.

Incorrect Answers:

Use a relational database when data that you work with is structured, and the structure is not subject to frequent changes.

Use Cloud storage (no relational database) for geographically distributed writes.

Reference:

<https://towardsdatascience.com/choosing-the-right-database-c45cd3a28f77>

 **Armandoo** Highly Voted 3 years, 3 months ago

Correct Answer: C (Transaction Writes)

upvoted 30 times

 **ranimswami** Highly Voted 3 years, 3 months ago

correct

upvoted 9 times

 **AGTraining** Most Recent 6 months, 2 weeks ago

Option C

upvoted 1 times

 **trashbox** 10 months, 2 weeks ago

C.

Appeared on 2023-11-05.

upvoted 1 times

 **MannyJZ** 1 year, 1 month ago

On exam 08/08/23

upvoted 1 times

 **VNR25** 1 year, 2 months ago

This question appeared in the 07/07/23 exam and the answer is correct.

upvoted 2 times

 **HRKK** 1 year, 2 months ago

CORRECT ANSWER C

upvoted 1 times

🗉 **PhyoPluto** 1 year, 6 months ago

the exam on 07/03/2023

upvoted 5 times

🗉 **chiiiweiii** 1 year, 6 months ago

Correct Answer.

upvoted 1 times

🗉 **King\_Lam** 1 year, 7 months ago

Was in exam 27/1/2023

upvoted 2 times

🗉 **Davisonye** 2 years ago

transactional writes. OLTP

upvoted 5 times

🗉 **Eltooth** 2 years, 2 months ago

Transactional writes

upvoted 3 times

🗉 **heenak55** 2 years, 4 months ago

In exam 21 April 2022

upvoted 4 times

🗉 **RichMelbourne** 2 years, 5 months ago

Shown on test 15/04/2022

upvoted 3 times

🗉 **AZ5cert** 2 years, 5 months ago

Correct Answer: C

upvoted 2 times

🗉 **RajeshAzure** 2 years, 6 months ago

RDBMS is for the transactional system so answer is correct.

upvoted 3 times

🗉 **sumanshu** 2 years, 7 months ago

Transactional Writes

upvoted 2 times

HOTSPOT -

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE: Each correct selection is worth one point.

Hot Area:

### Answer Area

Statements	Yes	No
Batch processing can output data to a file store	<input type="radio"/>	<input type="radio"/>
Batch processing can output data to a relational database	<input type="radio"/>	<input type="radio"/>
Batch processing can output data to a NoSQL database	<input type="radio"/>	<input type="radio"/>

Suggested Answer:

### Answer Area

Statements	Yes	No
Batch processing can output data to a file store	<input checked="" type="radio"/>	<input type="radio"/>
Batch processing can output data to a relational database	<input type="radio"/>	<input checked="" type="radio"/>
Batch processing can output data to a NoSQL database	<input type="radio"/>	<input checked="" type="radio"/>

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 -

Reference:

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

 **DemoBoy** Highly Voted 3 years, 5 months ago

All 3 are Yes. This is not just a big data question.

upvoted 237 times

 **Armandoo** Highly Voted 3 years, 3 months ago

YES, YES, YES

Batch processing can have output to SQL Database, Hive, HBase, Spark SQL. Look at the link and the diagram there

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

upvoted 67 times

 **K\_U\_99** Most Recent 1 month, 1 week ago

Does DP 900 have questions from these dumps only ? or should i read something else as well, please help

upvoted 1 times

 **Jeevi0905** 1 month, 3 weeks ago

Batch processing can output data to a file store:

Correct Answer: Yes

Explanation: Batch processing can indeed output data to a file store. This is commonly used for storing large volumes of data in files, such as logs, reports, or backups.

Batch processing can output data to a relational database:

Correct Answer: Yes

Explanation: Batch processing can also output data to a relational database. This is useful for scenarios where structured data needs to be stored and queried efficiently.

Batch processing can output data to a NoSQL database:

Correct Answer: Yes

Explanation: Yes, batch processing can also output data to a NoSQL database. NoSQL databases are designed to handle unstructured or semi-structured data, making them suitable for various use cases.

upvoted 2 times

  **AGTraining** 6 months, 2 weeks ago

YES

YES

YES

upvoted 1 times

  **Cabusas** 7 months, 1 week ago

I was thinking that all 3 are yes.

Based on Microsoft documentation there are 4 Sinks for stream processing namely Azure Event Hubs, Azure Data Lake Store Gen 2 or Azure blob storage, Azure SQL Database or Azure Synapse Analytics, or Azure Databricks, and Microsoft Power BI.

Which I think that Azure SQL Database is a relational.

Maybe the answer to this question is outdated based on the old technology.

upvoted 3 times

  **ravinperera** 9 months, 3 weeks ago

Statement 1: Batch processing can output data to a file store.

Yes. Batch processing can indeed output data to a file store. This is a common use case where processed data is stored in file systems like Azure Blob Storage or other similar cloud-based or on-premises file storage systems.

Statement 2: Batch processing can output data to a relational database.

Yes. Batch processing can output data to relational databases. After processing, the data can be structured and stored in relational database systems like Azure SQL Database, MySQL, PostgreSQL, or other relational databases.

Statement 3: Batch processing can output data to a NoSQL database.

Yes. Batch processing can also output data to NoSQL databases. This is particularly useful for unstructured or semi-structured data. NoSQL databases like Azure Cosmos DB, MongoDB, or others can be used for storing the results of batch processing tasks.

upvoted 5 times

  **Lscrinio** 10 months ago

YNN

"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."

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

upvoted 3 times

  **Law3** 10 months, 1 week ago

I have no experience with Azure batch, but seems it do not natively writes into sql, you need to add script and external libraries to run ETL operation within the batch job. Question is not clear at this point.

upvoted 2 times

  **MannyJZ** 1 year, 1 month ago

On exam 08/08/23

upvoted 1 times

🗨️ 👤 **akhil5432** 1 year, 1 month ago

yes

no

no

upvoted 1 times

🗨️ 👤 **Pistons** 1 year, 2 months ago

Frankly Speaking, In real work, we always have Batch Processing using Spark to write into Postgres...Perhaps it is not the 100% perfect architecture but it is used in many projets...And Theorically, it is possible although it is not really ETL (Data should be all transformed before loading )

upvoted 2 times

🗨️ 👤 **globy118** 1 year, 2 months ago

Was on exam 24/6/2023

upvoted 2 times

🗨️ 👤 **fafulka** 1 year, 3 months ago

in the exam 16/06/2023

upvoted 2 times

🗨️ 👤 **Dudures** 1 year, 5 months ago

Exam 04/15/2023

upvoted 6 times

🗨️ 👤 **XtraWest** 1 year, 5 months ago

Should be Y Y Y

upvoted 4 times

🗨️ 👤 **ahmedkmj** 1 year, 5 months ago

in the exam 26/03

upvoted 2 times

DRAG DROP -

Your company plans to load data from a customer relationship management (CRM) system to a data warehouse by using an extract, load, and transform (ELT) process.

Where does data processing occur for each stage of the ELT process? To answer, drag the appropriate locations to the correct stages. Each location 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:

**Locations**

- An in-memory data integration tool
- The CRM system
- The data warehouse

**Answer Area**

- Extract: Location
- Load: Location
- Transform: Location

**Suggested Answer:**

**Locations**

- An in-memory data integration tool
- The CRM system
- The data warehouse

**Answer Area**

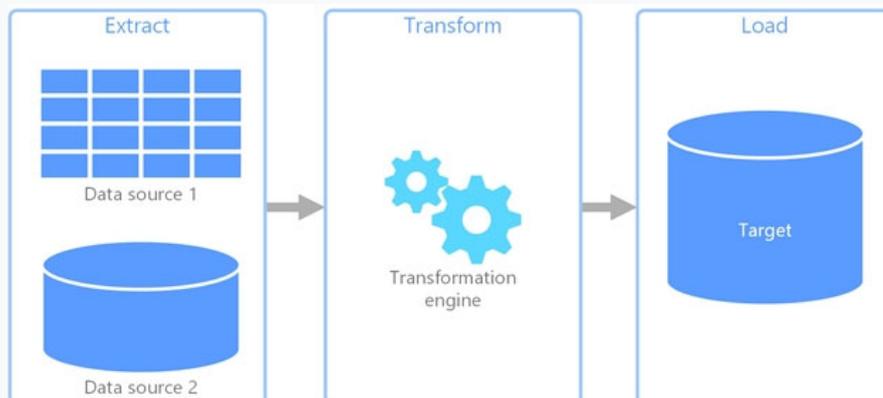
- Extract: The CRM system
- Load: The data warehouse
- Transform: An in-memory data integration tool

Box 1: The CRM system -  
Data is extracted from the CRM system.

Box 2: The data warehouse -  
Data is loaded to the data warehouse.

Box 3: An in-memory data integration tool

The data transformation that takes place usually involves various operations, such as filtering, sorting, aggregating, joining data, cleaning data, deduplicating, and validating data.



Reference:

<https://docs.microsoft.com/en-us/azure/architecture/data-guide/relational-data/etl>

The answer is CRM - DWH - DWH.

I got all the questions on analytics answered correctly in the exam.

upvoted 178 times

  **omw2wealth** 3 years, 1 month ago

Is it possible to drag the same option twice? :)

upvoted 4 times

  **researched\_answer\_boi** 3 years, 1 month ago

From the question: "Each location may be used once, more than once, or not at all.", so yep.

upvoted 13 times

  **[Removed]** 3 years, 3 months ago

I agree with your answer

upvoted 4 times

  **notem** 3 years, 2 months ago

Agree! Go to <https://docs.microsoft.com/en-us/azure/architecture/data-guide/relational-data/etl>

In the ELT pipeline, the transformation occurs in the target data store. Instead of using a separate transformation engine, the processing capabilities of the target data store are used to transform data.

upvoted 26 times

  **ANath** Highly Voted 3 years, 5 months ago

This question is related to ELT(Extract-Load-Transform) not ETL (Extract-Transform-Load). In ELT, transformation happens on Target. So, Load and Transform should happen on 'Data Warehouse'.

The answer given here is for ETL not for ELT.

upvoted 42 times

  **S320** 3 years, 3 months ago

I agree. The link to the MS site should be this sub-part further down the page:

<https://docs.microsoft.com/en-us/azure/architecture/data-guide/relational-data/etl#extract-load-and-transform-elt>

upvoted 8 times

  **AGTraining** Most Recent 6 months, 2 weeks ago

crm

datawarehouse

in memory data

upvoted 1 times

  **ff09375** 8 months, 1 week ago

FWIW, GPT 4 confirms

CRM

DWH

DWH

upvoted 1 times

  **sdg2844** 9 months, 1 week ago

I am wondering if there are two correct answers here, because with ELT, you actually extract > Load > Transform > Load that transformation somewhere else. So technically both DWH and in-memory data are theoretically correct.

upvoted 1 times

  **otapi** 7 months ago

The Transform can happen right in the DWH, why would you load the data 2nd time to somewhere else? It is possible to add as many extra load and transform as you wish, but that will not be an efficient ELT.

upvoted 2 times

  **MannyJZ** 1 year, 1 month ago

On exam 08/08/23

upvoted 2 times

  **akhil5432** 1 year, 1 month ago

crm

datawarehouse

in memory data

upvoted 1 times

🗨️ **fafulka** 1 year, 3 months ago

in the exam 16/06/2023

upvoted 1 times

🗨️ **Maddy4** 1 year, 2 months ago

what was your answer? CRM - DWH - DWH?

upvoted 3 times

🗨️ **Dudures** 1 year, 5 months ago

Exam 04/15/2023

upvoted 4 times

🗨️ **TestTurtle468** 1 year, 5 months ago

On the following website: <https://learn.microsoft.com/en-us/azure/architecture/data-guide/relational-data/etl>, the diagram in the result is an example of ETL, but the question discusses ELT.

Is this correct, or am I thinking too much on this question? Please advise and thanks.

upvoted 1 times

🗨️ **chiiweiii** 1 year, 6 months ago

Correct Answer.

upvoted 1 times

🗨️ **Joanna0** 1 year, 7 months ago

This is ELT.

CRM - EDW - EDW is the correct answer.

upvoted 3 times

🗨️ **edurakhan** 1 year, 10 months ago

Exam 11/15/22

upvoted 2 times

🗨️ **Tonnie85** 1 year, 10 months ago

For ELT, the correct answer is CRM-DWH-DWH

<https://learn.microsoft.com/en-us/azure/architecture/data-guide/relational-data/etl>

Extract, load, and transform (ELT) differs from ETL solely in where the transformation takes place. In the ELT pipeline, the transformation occurs in the target data store. Instead of using a separate transformation engine, the processing capabilities of the target data store are used to transform data.

upvoted 3 times

🗨️ **ThomasDB** 1 year, 11 months ago

Exam: 12/10/2022

upvoted 2 times

🗨️ **delmarino** 1 year, 11 months ago

The given answer is correct

CRM

DWH

IN memory data

upvoted 3 times

🗨️ **DUTI** 2 years ago

ELT is different from ETL and picture in explanation is wrong... It is CRM - DWH - DWH.

upvoted 1 times

HOTSPOT -

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

Hot Area:

### Answer Area

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

▼
cognitive
descriptive
predictive
prescriptive

Suggested Answer:

### Answer Area

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

▼
cognitive
descriptive
predictive
prescriptive

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?

Reference:

<https://azure.microsoft.com/en-us/blog/answering-whats-happening-whys-happening-and-what-will-happen-with-iot-analytics/>

 **hmAZtime** Highly Voted 3 years, 5 months ago

Descriptive

upvoted 65 times

 **ANath** Highly Voted 3 years, 5 months ago

Descriptive is 100% correct.

upvoted 20 times

 **AGTraining** Most Recent 6 months, 2 weeks ago

Descriptive

upvoted 1 times

 **Lscranio** 10 months ago

Descriptive.

upvoted 1 times

 **akhil5432** 1 year, 1 month ago

descriptive

upvoted 2 times

 **chiiiweiii** 1 year, 6 months ago

Correct Answer.

upvoted 2 times

 **Joanna0** 1 year, 7 months ago

Descriptive!

upvoted 1 times

🗨️ 👤 **shubham270319** 1 year, 8 months ago

Descriptive Analytics

upvoted 1 times

🗨️ 👤 **Sush1994** 1 year, 8 months ago

Descriptive

upvoted 1 times

🗨️ 👤 **Edriv** 1 year, 9 months ago

Cognitive

upvoted 1 times

🗨️ 👤 **Eltooth** 2 years, 2 months ago

Descriptive

upvoted 2 times

🗨️ 👤 **Kiki\_25** 2 years, 3 months ago

definetely descriptive

upvoted 3 times

🗨️ 👤 **megalai** 2 years, 4 months ago

descriptive

upvoted 2 times

🗨️ 👤 **cormorant** 2 years, 5 months ago

telling it like it is: descriptive

upvoted 2 times

🗨️ 👤 **AZ5cert** 2 years, 5 months ago

Descriptive

upvoted 1 times

🗨️ 👤 **RajeshAzure** 2 years, 6 months ago

Descriptive

upvoted 1 times

🗨️ 👤 **sumanshu** 2 years, 7 months ago

Descriptive.

upvoted 1 times

DRAG DROP -

Match the types of visualizations to the appropriate descriptions.

To answer, drag the appropriate visualization type from the column on the left to its description on the right. Each visualization type may be used once, more than once, or not at all.

NOTE: Each correct match is worth one point.

Select and Place:

#### Visualizations

Key influencer

Scatter

Treemap

#### Answer Area

Visualization

A chart of colored, nested rectangles that displays individual data points represented by the size and color of a relative rectangle.

Visualization

A chart that displays the major contributors of a selected result or value.

Visualization

A chart that shows the relationship between two numerical values.

#### Suggested Answer:

#### Visualizations

Key influencer

Scatter

Treemap

#### Answer Area

Treemap

A chart of colored, nested rectangles that displays individual data points represented by the size and color of a relative rectangle.

Key influencer

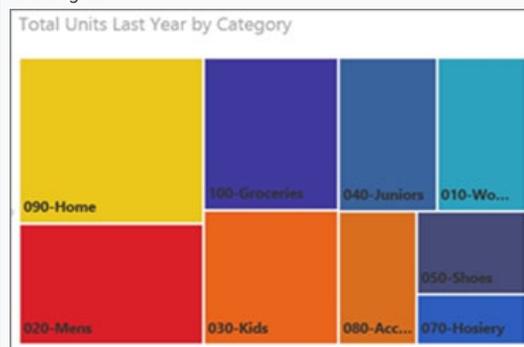
A chart that displays the major contributors of a selected result or value.

Scatter

A chart that shows the relationship between two numerical values.

Box 1: Tree map -

Treemaps are charts of colored rectangles, with size representing value. They can be hierarchical, with rectangles nested within the main rectangles.



Box 2: Key influencer -

A key influencer chart displays the major contributors to a selected result or value.

Box 3: Scatter -

Scatter and Bubble charts display relationships between 2 (scatter) or 3 (bubble) quantitative measures -- whether or not, in which order, etc.

**Crazy\_Angel** Highly Voted 3 years, 2 months ago

Treemap : Are charts of colored rectangles, with size representing value. They can be hierarchical, with rectangles nested within the main rectangles.

Key Influencer : A key influencer chart displays the major contributors to a selected result or value.

scatter : Display relationships between 2 (scatter) or 3 (bubble) quantitative measures -- whether or not, in which order, etc.  
upvoted 55 times

**yeanlingmedal71** Highly Voted 3 years, 3 months ago

upvoted 26 times

  **AVP\_Riga** 3 years, 3 months ago

Thanks!

upvoted 1 times

  **AGTraining** Most Recent 6 months, 2 weeks ago

treemap

keyinfluencer

scatter

upvoted 1 times

  **akhil5432** 1 year, 1 month ago

treemap

keyinfluencer

scatter

upvoted 1 times

  **mamahani** 1 year, 5 months ago

I took the exam 12.04.2023 and got 1000, 100% score; this question appeared, correct answer:

treemap - key influences - scatter

upvoted 5 times

  **chiiweiii** 1 year, 6 months ago

Correct Answer.

upvoted 2 times

  **Joanna0** 1 year, 7 months ago

The correct answer:

Treemap

Key influencer

Scatter

upvoted 2 times

  **Drchatss** 2 years, 1 month ago

01/08.....EXAM

upvoted 2 times

  **Eltooth** 2 years, 2 months ago

Tree

Key

Scatter

upvoted 4 times

  **dk87** 2 years, 4 months ago

Exam: 05/18/2022

upvoted 2 times

  **mre** 2 years, 5 months ago

Correct treemap, key inf, scatter

upvoted 4 times

  **AZ5cert** 2 years, 5 months ago

Treemap : Are charts of colored rectangles, with size representing value. They can be hierarchical, with rectangles nested within the main rectangles.

Key Influencer : A key influencer chart displays the major contributors to a selected result or value.

scatter : Display relationships between 2 (scatter) or 3 (bubble) quantitative measures -- whether or not, in which order, etc.

upvoted 3 times

  **Shuwy** 2 years, 6 months ago

Correct!

upvoted 3 times

  **Moezey** 2 years, 7 months ago

CORRECT

upvoted 3 times

🗨️ 👤 **sumanshu** 2 years, 7 months ago

Given ans is correct

upvoted 1 times

🗨️ 👤 **Mjwc** 2 years, 11 months ago

Correct according to Microsoft Learning.

Source:

<https://docs.microsoft.com/en-us/learn/modules/explore-concepts-of-data-analytics/3-explore-data-visualization>

upvoted 5 times

🗨️ 👤 **Mishra\_ji** 2 years, 12 months ago

This is correct

upvoted 2 times

You need to create an Azure Storage account.

Data in the account must replicate outside the Azure region automatically.

Which two types of replication can you use for the storage account? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

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

**Suggested Answer: BD**

D: Azure Storage offers two options for copying your data to a secondary region:

- ⇒ Geo-redundant storage (GRS)
- ⇒ Geo-zone-redundant storage (GZRS)

B: 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).

Reference:

<https://docs.microsoft.com/en-us/azure/storage/common/storage-redundancy#redundancy-in-a-secondary-region>

Community vote distribution



🗳️ **Feyisara** Highly Voted 3 years, 5 months ago

The chosen option is correct given the requirements in the question.

upvoted 47 times

🗳️ **AlfuryDB** Highly Voted 3 years, 6 months ago

Careful with this one. If we are talking about "outside replication", the options will be GRS and RA-GRS (particular case that gives active-active between regions with reading accessible in the second one). However, GRS does NOT allow to read in the secondary region in an active-active way. It only happens when the first region fails. Bear this in mind. The option is correct given the requirements in the question.

upvoted 38 times

🗳️ **GuruvuManoj** Most Recent 5 months, 2 weeks ago

Selected Answer: D

In most of the cases, GRS would be preferred, so few of the business would request for access for writing like to create a folder inside blob, or to upload a doc manually into the blob, so most of the cases GRS is preferable.

upvoted 1 times

🗳️ **AGTraining** 6 months, 2 weeks ago

Option B and C

upvoted 1 times

🗳️ **Nicolade** 8 months, 2 weeks ago

B and D are correct

upvoted 1 times

🗳️ **gcsilva83** 9 months, 2 weeks ago

Selected Answer: BD

B AND D

upvoted 1 times

🗳️ **hcq31818** 11 months, 4 weeks ago

Selected Answer: BD

B and D

upvoted 1 times

🗨️ **akhil5432** 1 year, 1 month ago

B AND D

upvoted 2 times

🗨️ **rohirag** 1 year, 2 months ago

**Selected Answer: BD**

Answer

upvoted 2 times

🗨️ **rohirag** 1 year, 2 months ago

BD are the correct answers.

upvoted 1 times

🗨️ **HRKK** 1 year, 2 months ago

B AND D IS THE CORRECT ANSWER

upvoted 1 times

🗨️ **Css6000** 1 year, 5 months ago

B and D are correct.

upvoted 2 times

🗨️ **Steph2000** 1 year, 6 months ago

**Selected Answer: BD**

Both are a secondary location outside the region

upvoted 2 times

🗨️ **chiiiweiii** 1 year, 6 months ago

**Selected Answer: BD**

Correct Answer.

upvoted 1 times

🗨️ **Itsplata** 1 year, 6 months ago

**Selected Answer: BD**

Correct

upvoted 1 times

🗨️ **MallehJolad** 1 year, 7 months ago

B and D are the correct answers

upvoted 1 times

🗨️ **raydel92** 1 year, 11 months ago

**Selected Answer: BD**

B and D

upvoted 5 times

HOTSPOT -

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE: Each correct selection is worth one point.

Hot Area:

### Answer Area

Statements	Yes	No
Platform as a service (PaaS) database offerings in Azure require less setup and configuration effort than infrastructure as a service (IaaS) database offerings.	<input type="radio"/>	<input type="radio"/>
Platform as a service (PaaS) database offerings in Azure provide end users with the ability to control and update the operating system version.	<input type="radio"/>	<input type="radio"/>
All relation and non-relational platform as a service (PaaS) database offerings in Azure can be paused to reduce costs.	<input type="radio"/>	<input type="radio"/>

Suggested Answer:

### Answer Area

Statements	Yes	No
Platform as a service (PaaS) database offerings in Azure require less setup and configuration effort than infrastructure as a service (IaaS) database offerings.	<input checked="" type="radio"/>	<input type="radio"/>
Platform as a service (PaaS) database offerings in Azure provide end users with the ability to control and update the operating system version.	<input type="radio"/>	<input checked="" type="radio"/>
All relation and non-relational platform as a service (PaaS) database offerings in Azure can be paused to reduce costs.	<input type="radio"/>	<input checked="" type="radio"/>

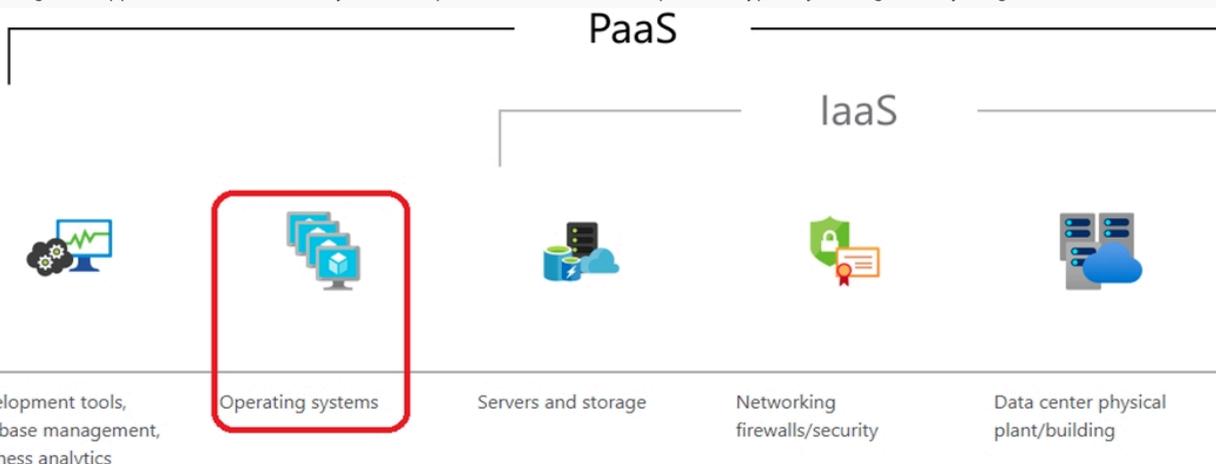
Box 1: Yes -

Like IaaS, PaaS includes infrastructure " servers, storage, and networking " but also middleware, development tools, business intelligence (BI) services, database management systems, and more. PaaS is designed to support the complete web application lifecycle: building, testing, deploying, managing, and updating.

PaaS allows you to avoid the expense and complexity of buying and managing software licenses, the underlying application infrastructure and middleware, container orchestrators such as Kubernetes, or the development tools and other resources

Box 2: No -

You manage the applications and services you develop, and the cloud service provider typically manages everything else.



Box 3: No -

There really is no way to pause / stop billing for your Azure SQL Database.

Microsoft's official answer "Yes, you can export your database. Delete the Azure SQL database and that will pause billing. Then when you need it you can create a new database and import your previously exported DB."

Reference:

<https://azure.microsoft.com/en-us/overview/what-is-paas>

  **sakanan** Highly Voted 2 years, 4 months ago

correct

upvoted 35 times

  **Edriv** Highly Voted 2 years, 2 months ago

YES, NO, YES

upvoted 14 times

  **Jezza8024** 2 months, 1 week ago

Yeah dont be an idiota

upvoted 1 times

  **HananS** 2 years, 1 month ago

Don't be an IDIO.

upvoted 18 times

  **Jeevi0905** Most Recent 1 month, 3 weeks ago

Yes,NO,Yes

upvoted 1 times

  **jerrybin30731** 2 months, 2 weeks ago

what is the "end users" mean? developer? or is like player or something like that?

upvoted 1 times

  **Ruesam** 4 months ago

Yes, No, No.

I see there are a lot of questions about the last one. The answer is no because it's true that at the moment, you cannot pause Azure SQL DB:

<https://techcommunity.microsoft.com/t5/azure-database-support-blog/stop-azure-sql-database/ba-p/3890369#:~:text=Stopping%20an%20Azure%20SQL%20Database,SQL%20Database%20serverless%20compute%20tier.>

upvoted 1 times

  **AGTraining** 6 months, 2 weeks ago

YES, NO, NO

upvoted 1 times

  **Cabusas** 7 months, 1 week ago

For me it is YES, NO, YES.

Because in an Azure SQL Managed Instance you can pause it.

Refer: <https://learn.microsoft.com/en-us/azure/azure-sql/managed-instance/instance-stop-start-how-to?view=azuresql&tabs=azure-powershell-prep%2Cazure-portal>

upvoted 1 times

  **trashbox** 10 months, 2 weeks ago

Yes, No, No.

Appeared on 2023-11-05.

upvoted 2 times

  **DonDemik** 1 year ago

Last question is tricky : not all databases can be paused. f.x. Cosmos DB. However, some services such as Azure SQL database provides an auto-pause feature that user can configure.

[https://learn.microsoft.com/en-us/answers/questions/707576/how-do-i-pause-an-azure-sql-database-\(serverless\)](https://learn.microsoft.com/en-us/answers/questions/707576/how-do-i-pause-an-azure-sql-database-(serverless))

upvoted 3 times

  **FedeMI** 1 year ago

Corret way is: YES, NO, NO

upvoted 3 times

🗨️ **MannyJZ** 1 year, 1 month ago

On exam today 08/08/23

upvoted 1 times

🗨️ **akhil5432** 1 year, 1 month ago

YES

NO

NO

upvoted 2 times

🗨️ **VNR25** 1 year, 2 months ago

This question appeared in the 07/07/23 exam and the answer is correct.

upvoted 5 times

🗨️ **XtraWest** 1 year, 6 months ago

YES, NO, NO (seems correct)

upvoted 4 times

🗨️ **chiiiweiii** 1 year, 6 months ago

Correct Answer.

upvoted 2 times

🗨️ **Sand9** 1 year, 8 months ago

YES, NO, NO

upvoted 6 times

🗨️ **Wobbla** 1 year, 9 months ago

ignore this, it asks for all relational and non-relational PaaS offerings

upvoted 2 times

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

- A. REVOKE
- B. DISABLE
- C. INSERT
- D. GRANT

**Suggested Answer:** C

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

Reference:

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

*Community vote distribution*

C (100%)

🗉 **tousdev** 5 months ago

**Selected Answer: C**

DML : Insert / Update / Delete  
upvoted 1 times

🗉 **GuruvuManoj** 5 months, 2 weeks ago

**Selected Answer: C**

"INSERT" is one of the Data Manipulation function as we can add multiple lines of rows into a table.  
upvoted 2 times

🗉 **AGTraining** 6 months, 2 weeks ago

OPTION C

upvoted 1 times

🗉 **Nicolade** 8 months, 2 weeks ago

The answer is C

upvoted 1 times

🗉 **kjszala** 11 months, 3 weeks ago

**Selected Answer: C**

Insert is correct  
upvoted 1 times

🗉 **hcq31818** 1 year ago

**Selected Answer: C**

Insert is DML  
upvoted 1 times

🗉 **akhil5432** 1 year, 1 month ago

CORRECT OPTION IS C

upvoted 2 times

🗉 **rohirag** 1 year, 2 months ago

**Selected Answer: C**

C is correct. CRUD Operations are part of DML.

upvoted 2 times

🗨️ **Css6000** 1 year, 5 months ago

DCL includes commands such as GRANT and REVOKE. TCL commands include BEGIN, COMMIT, ROLLBACK, SAVEPOINT, SET TRANSACTION

upvoted 1 times

🗨️ **Varuntv** 1 year, 5 months ago

**Selected Answer: C**

insert is the correct answer

upvoted 3 times

🗨️ **chiiiweiii** 1 year, 6 months ago

**Selected Answer: C**

Correct Answer.

upvoted 4 times

🗨️ **BaltVargas** 1 year, 6 months ago

insert: DML (Data Manipulation Language)

GRANT, DENY, REVOKE: DCL (Data Control Language)

upvoted 1 times

🗨️ **smgk** 1 year, 7 months ago

insert

upvoted 1 times

🗨️ **gdwlnaveen** 1 year, 7 months ago

insert is correct answer

upvoted 1 times

🗨️ **Anil\_Mahanty** 1 year, 8 months ago

Insert is the correct answer

upvoted 2 times

🗨️ **yenada** 1 year, 8 months ago

**Selected Answer: C**

Insert is the correct answer

upvoted 1 times

🗨️ **aPinkFluffyYak** 1 year, 9 months ago

**Selected Answer: C**

Insert is the correct answer

upvoted 1 times

You have a SQL query that combines customer data and order data. The query includes calculated columns. You need to create a database object that would allow other users to rerun the same SQL query. What should you create?

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

**Suggested Answer: B**

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.

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/views/views>

Community vote distribution

B (100%)

 **puksometron** Highly Voted 2 years, 11 months ago

Answer is correct  
upvoted 29 times

 **heenak55** Highly Voted 2 years, 4 months ago

In exam 21 April 2022  
upvoted 5 times

 **tousdev** Most Recent 5 months ago

Selected Answer: B

Answer B. View can be used by other users  
upvoted 1 times

 **AGTraining** 6 months, 2 weeks ago

Option B  
upvoted 1 times

 **trashbox** 10 months, 2 weeks ago

Selected Answer: B

Appeared on 2023-11-05.  
upvoted 1 times

 **kjszala** 11 months, 3 weeks ago

Selected Answer: B

Correct  
upvoted 3 times

 **TKC\_2023** 1 year ago

Selected Answer: B

correct  
upvoted 1 times

 **hcq31818** 1 year ago

Selected Answer: B

View is the correct answer  
upvoted 1 times

 **AZFabio** 1 year ago

Selected Answer: B

A view allows you to reuse a query  
upvoted 1 times

- 🗨️ **MannyJZ** 1 year, 1 month ago  
On exam today 08/08/23  
upvoted 1 times
- 🗨️ **akhil5432** 1 year, 1 month ago  
OPTION B  
upvoted 1 times
- 🗨️ **rohirag** 1 year, 2 months ago  
**Selected Answer: B**  
Create a VIEW  
upvoted 1 times
- 🗨️ **VNR25** 1 year, 2 months ago  
This question appeared in the 07/07/23 exam and the answer is correct.  
upvoted 1 times
- 🗨️ **vicks60** 1 year, 5 months ago  
**Selected Answer: B**  
View is correct answer  
upvoted 2 times
- 🗨️ **chiiiweiii** 1 year, 6 months ago  
**Selected Answer: B**  
Correct Answer.  
upvoted 2 times
- 🗨️ **MalleshJolad** 1 year, 7 months ago  
Option B  
upvoted 1 times
- 🗨️ **prav936** 1 year, 10 months ago  
answer is correct  
upvoted 3 times

HOTSPOT -

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

Hot Area:

### Answer Area

A key/value data store is optimized for

	▼
enforcing constraints	
simple lookups	
table joins	
transactions	

### Answer Area

Suggested Answer:

A key/value data store is optimized for

	▼
enforcing constraints	
simple lookups	
table joins	
transactions	

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.

Reference:

<https://docs.microsoft.com/en-us/azure/architecture/guide/technology-choices/data-store-overview>

- 🗨️ **hmAZtime** Highly Voted 3 years, 5 months ago  
 right answer  
 upvoted 43 times
- 🗨️ **RichMelbourne** Highly Voted 2 years, 5 months ago  
 Shown on test 15/04/2022  
 upvoted 12 times
- 🗨️ **AGTraining** Most Recent 6 months, 2 weeks ago  
 Simple lookups  
 upvoted 1 times
- 🗨️ **holymolly** 7 months, 2 weeks ago  
 contact to get full questions  
 upvoted 1 times
- 🗨️ **Crocodile181** 1 month, 4 weeks ago  
 me want  
 upvoted 1 times
- 🗨️ **james2033** 10 months, 3 weeks ago  
 A key/value data store is optimized for SIMPLE LOOKUPS.  
 upvoted 1 times
- 🗨️ **Han\_** 1 year ago  
 Simple lookups  
 upvoted 2 times

- 🗨️ **MannyJZ** 1 year, 1 month ago  
On exam today 08/08/23  
upvoted 1 times
- 🗨️ **akhil5432** 1 year, 1 month ago  
SIMPLE LOOKUPS  
upvoted 1 times
- 🗨️ **zanjeev** 1 year, 2 months ago  
Simple Lookup  
upvoted 1 times
- 🗨️ **fafulka** 1 year, 3 months ago  
Exam 16/06/2023  
upvoted 1 times
- 🗨️ **98761234** 1 year, 5 months ago  
Shown in exam 04/2023  
upvoted 4 times
- 🗨️ **Dudures** 1 year, 5 months ago  
Exam 04/15/2023  
upvoted 2 times
- 🗨️ **mamahani** 1 year, 5 months ago  
I took the exam 12.04.2023 and got 1000, 100% score; this question appeared, correct answer:  
simple lookups  
upvoted 6 times
- 🗨️ **ahmedkmj** 1 year, 5 months ago  
in the exam 26/03  
upvoted 1 times
- 🗨️ **chiiweiii** 1 year, 6 months ago  
Correct Answer.  
upvoted 1 times
- 🗨️ **edurakhan** 1 year, 10 months ago  
Exam 11/15/22  
upvoted 1 times
- 🗨️ **ThomasDB** 1 year, 11 months ago  
Exam: 12/10/2022  
upvoted 3 times

DRAG DROP -

Match the types of data to the appropriate Azure data services.

To answer, drag the appropriate data type from the column on the left to its service on the right. Each data type may be used once, more than once, or not at all.

NOTE: Each correct match is worth one point.

Select and Place:

Data Types	Answer Area
Image files	Data type Azure Blob storage
Key/value pairs	Data type Azure Cosmos DB Gremlin API
Relationships between employees	Data type Azure Table storage

Suggested Answer:

Data Types	Answer Area
Image files	Image files Azure Blob storage
Key/value pairs	Key/value pairs Azure Cosmos DB Gremlin API
Relationships between employees	Relationships between employees Azure Table storage

Box 1: Image files -

Azure Blob storage is suitable for image files.

Box 2:Key/value pairs -

Azure CosmosDB table API is a key-value storage hosted in the cloud.

Box 3: Relationship between employees

One-to-many relationships between business domain objects occur frequently: for example, one department has many employees. There are several ways to implement one-to-many relationships in the Azure Table service.

Reference:

<https://docs.microsoft.com/en-us/azure/storage/tables/table-storage-design-modeling>

 **tibi1233**  3 years, 5 months ago

1. Image files = Azure Blob storage
  2. Relationship between employees = Azure Cosmos DB Gremlin API
  3. Key/value pairs = Azure Table Storage
- upvoted 778 times

 **Tsbki** 1 month ago  
767 upvotes and still no correction  
upvoted 1 times

 **MPB** 2 years, 9 months ago  
correct  
upvoted 2 times

 **TJ001** 2 years, 8 months ago  
agree with the answer  
upvoted 3 times

🗨️ 👤 **Iroshan** 3 years, 5 months ago

Agreed.

This is the correct answer.

upvoted 3 times

🗨️ 👤 **no121251** Highly Voted 3 years, 1 month ago

Did the person who prepared these answers even pass the test themselves? This isn't even an accidental swap, just a clear misunderstanding from the person who wrote the answer. The Azure Table storage module on MS Learn EXPLICITLY STATES that a weakness of Table storage is its inability to document relationships between rows/entities in the table.

upvoted 87 times

🗨️ 👤 **AGTraining** Most Recent 6 months, 2 weeks ago

1. Image files = Azure Blob storage
2. Relationship between employyes = Azure Cosmos DB Gremlin API
3. Key/value pairs = Azure Table Storage

upvoted 2 times

🗨️ 👤 **jxs221** 7 months ago

Image files -> Azure Blob storage

Key/value pairs -> Azure Table storage

Relationships between employees -> Azure Cosmos DB Gremlin API

upvoted 2 times

🗨️ 👤 **Rudramina** 9 months, 3 weeks ago

Image files = Azure Blob storage

Relationship between employees = Azure Cosmos DB Gremlin API

Key/value pairs = Azure Table Storage

upvoted 6 times

🗨️ 👤 **masadamsahid** 10 months, 3 weeks ago

Should be:

- Image File = Blob

- Relationship between employees = CosmosDB (Gremlin API)

- Key/value pairs = AZ Table Storage

upvoted 4 times

🗨️ 👤 **james2033** 10 months, 3 weeks ago

'Image files' --> 'Azure Blob storage'

'Key/value pairs' --> 'Azure Cosmos DB Gremlin API'

'Relationships between employees' --> 'Azure Table storage'

upvoted 1 times

🗨️ 👤 **TKC\_2023** 1 year ago

1. Image files = Azure Blob storage

2. Relationship between employyes = Azure Cosmos DB Gremlin API

3. Key/value pairs = Azure Table Storage

upvoted 3 times

🗨️ 👤 **MannyJZ** 1 year, 1 month ago

On exam today 08/08/23

upvoted 2 times

🗨️ 👤 **akhil5432** 1 year, 1 month ago

SAME sequence in answer row

upvoted 1 times

🗨️ 👤 **Pady1234** 1 year, 2 months ago

1. Image files = Azure Blob storage

2. Relationship between employees = Azure Cosmos DB Gremlin API

3. Key/value pairs = Azure Table Storage

upvoted 1 times

🗨️ 👤 **HRKK** 1 year, 2 months ago

1. Image files = Azure Blob storage
2. Relationship between employees = Azure Cosmos DB Gremlin API
3. Key/value pairs = Azure Table Storage

upvoted 2 times

🗨️ 👤 **fafulka** 1 year, 3 months ago

Exam 16/06/2023

upvoted 3 times

🗨️ 👤 **Elderbds** 1 year, 3 months ago

De acordo com o BING AI: Os tipos de dados podem ser combinados com os serviços de dados do Azure da seguinte maneira:

Arquivos de imagem podem ser armazenados no Armazenamento de Blobs do Azure, que é um serviço de armazenamento de objetos para armazenar grandes quantidades de dados não estruturados.

Pares chave/valor podem ser armazenados no Armazenamento de Tabelas do Azure, que é um serviço NoSQL que armazena dados estruturados em tabelas particionadas.

Relações entre funcionários podem ser representadas usando o API Gremlin do Azure Cosmos DB, que é um serviço de banco de dados gráfico totalmente gerenciado que permite modelar e consultar dados relacionais complexos.

upvoted 4 times

🗨️ 👤 **Vanessa23** 1 year, 5 months ago

Relationships = Gremlin (nodes, edges)

Image = Blob

Table = Key/values

upvoted 4 times

🗨️ 👤 **Dudures** 1 year, 5 months ago

Exam 04/15/2023

upvoted 5 times

🗨️ 👤 **A\_Fardeen** 1 year, 6 months ago

1. Image files = Azure Blob storage

2. Relationship between employees = Azure Cosmos DB Gremlin API

3. Key/value pairs = Azure Table Storage

upvoted 3 times

DRAG DROP -

Match the Azure Data Lake Storage Gen2 terms to the appropriate levels in the hierarchy.

To answer, drag the appropriate term from the column on the left to its level on the right. Each term may be used once, more than once, or not at all.

NOTE: Each correct match is worth one point.

Select and Place:

**Terms**

Azure Storage account

File share

Container

**Answer Area**

Azure Resource Group

Term

Term

Folders

Files

**Suggested Answer:**

**Terms**

Azure Storage account

File share

Container

**Answer Area**

Azure Resource Group

Azure Storage account

File share

Folders

Files

Box 1: Azure Storage account -

Azure file shares are deployed into storage accounts, which are top-level objects that represent a shared pool of storage.

Box 2: File share -

Reference:

<https://docs.microsoft.com/en-us/azure/storage/files/storage-how-to-create-file-share>

**seantb** Highly Voted 3 years, 5 months ago

The question says this is an ADL, in which case you would create a container not a file share... you can have folders and files in the container  
upvoted 285 times

**niw66966** 2 years, 9 months ago

Yep, agree.

upvoted 3 times

**GP2000** 2 years, 4 months ago

Agree with your answer

upvoted 2 times

**wellingtonluis** 2 years, 10 months ago

Absolutely agree

upvoted 1 times

- 🗄️ 👤 **rajat7878** 2 years, 6 months ago  
yes, agreed  
upvoted 1 times
- 🗄️ 👤 **wejapec21** Highly Voted 🏆 3 years, 5 months ago  
Azure storage account and container make more sense.  
upvoted 130 times
- 🗄️ 👤 **TKC\_2023** Most Recent 🕒 1 month, 1 week ago  
Azure Storage Account >> Container  
upvoted 1 times
- 🗄️ 👤 **Sr18** 4 months, 2 weeks ago  
Blob Storage / Container / folders.../ Files  
upvoted 2 times
- 🗄️ 👤 **QeUtPat** 5 months, 2 weeks ago  
Azure storage account and container are the correct answers.  
upvoted 1 times
- 🗄️ 👤 **AGTraining** 6 months, 2 weeks ago  
Azure Storage Account >> Container  
upvoted 2 times
- 🗄️ 👤 **Yomzie** 6 months, 3 weeks ago  
The correct sequence for ADLS Gen2:  
Azure Storage Account >> Container >> Directories >> Files

Azure Data Lake Storage Gen2 builds on blob storage and optimizes I/O of high-volume data by using a hierarchical namespace that organizes blob data into directories, and stores metadata about each directory and the files within it. This structure allows operations, such as directory renames and deletes, to be performed in a single atomic operation. Flat namespaces, by contrast, require several operations proportionate to the number of objects in the structure. Hierarchical namespaces keep the data organized, which yields better storage and retrieval performance for an analytical use case and lowers the cost of analysis.

SOURCE: <https://learn.microsoft.com/en-us/training/modules/introduction-to-azure-data-lake-storage/4-azure-data-lake-and-blob-storage>  
upvoted 1 times

- 🗄️ 👤 **Sayuri1313** 9 months, 2 weeks ago  
Account -> Container  
  
<https://learn.microsoft.com/en-us/training/modules/introduction-to-azure-data-lake-storage/4-azure-data-lake-and-blob-storage>  
upvoted 4 times
- 🗄️ 👤 **james2033** 10 months, 3 weeks ago  
Azure Resource Group --> Azure Storage account --> File share --> Folders --> Files.  
  
<https://learn.microsoft.com/en-us/azure/storage/files/storage-how-to-create-file-share?tabs=azure-portal#change-the-tier-of-an-azure-file-share>  
upvoted 1 times
- 🗄️ 👤 **otapi** 7 months ago  
The question was on context of Azure Data Lake Storage Gen2. File share is a different service.  
upvoted 1 times
- 🗄️ 👤 **Nino\_Sasania** 11 months, 1 week ago  
This question was on the exam today 12/10/23  
upvoted 3 times
- 🗄️ 👤 **TKC\_2023** 1 year ago  
Azure storage account and container  
upvoted 1 times
- 🗄️ 👤 **hcq31818** 1 year ago  
In azure data lake storage gen2, the hierarchy level is Azure subscription --> Az. Resource grps --> Az. data lake storage gen2 account--> Containers-->Folders & Files

upvoted 2 times

🗨️ 👤 **akhil5432** 1 year, 1 month ago  
azure storage account and file share  
upvoted 1 times

🗨️ 👤 **MikeScout** 1 year, 3 months ago  
Create a Datalake Gen2 in Azure. 1st a Storage account and 2nd check the change of the container icon. Open it and download the files. Hence, it is Storage account and Container.  
upvoted 1 times

🗨️ 👤 **rgmar** 1 year, 3 months ago  
Answer should be Azure Storage account and container  
upvoted 4 times

🗨️ 👤 **DataEngDP** 1 year, 4 months ago  
Correct answer is Container, but there is a difference in terminology stated by Azure. Container and File systems are synonymous. See:  
<https://learn.microsoft.com/en-us/azure/storage/blobs/data-lake-storage-introduction>  
upvoted 1 times

🗨️ 👤 **masadamsahid** 10 months, 3 weeks ago  
In the question above, it's written as "File Share" not "File System". So, I don't think we need to be confused about the synonym.  
upvoted 1 times

🗨️ 👤 **mamahani** 1 year, 5 months ago  
I took the exam 12.04.2023 and got 1000, 100% score; this question appeared, correct answer:  
Azure Storage account - Container  
upvoted 24 times

What are two characteristics of real-time data processing? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. Data is processed periodically
- B. Low latency is expected
- C. High latency is acceptable
- D. Data is processed as it is created

**Suggested Answer:** BD

Real time processing deals with streams of data that are captured in real-time and processed with minimal latency to generate real-time (or near-real-time) reports or automated responses.

Reference:

<https://docs.microsoft.com/en-us/azure/architecture/data-guide/big-data/real-time-processing>

Community vote distribution

BD (100%)

- 🗲 **Nirav2387** Highly Voted 3 years, 2 months ago  
B&D is the answer  
upvoted 61 times
- 🗲 **ANath** Highly Voted 3 years, 5 months ago  
Correct  
upvoted 11 times
- 🗲 **profesorklaus** Most Recent 1 week, 3 days ago  
Selected Answer: BD  
Correct  
upvoted 1 times
- 🗲 **3jess** 2 months ago  
Selected Answer: BD  
Correct  
upvoted 1 times
- 🗲 **tousdev** 5 months ago  
Selected Answer: BD  
B&D answers  
upvoted 2 times
- 🗲 **okayhey** 6 months, 1 week ago  
BD  
would not be realtime if you have high latency, now would it ;)  
upvoted 1 times
- 🗲 **AGTraining** 6 months, 2 weeks ago  
option B e D  
upvoted 1 times
- 🗲 **Bissto** 11 months, 2 weeks ago  
Selected Answer: BD  
correct  
upvoted 2 times
- 🗲 **kjszala** 11 months, 3 weeks ago  
Selected Answer: BD  
correct  
upvoted 1 times

TKC\_2023 1 year ago

Selected Answer: BD

B, D are correct  
upvoted 1 times

rohirag 1 year, 2 months ago

Selected Answer: BD

Correct answers.  
upvoted 1 times

vicks60 1 year, 5 months ago

Selected Answer: BD

B and D is correct Answer  
upvoted 2 times

Varuntv 1 year, 5 months ago

Selected Answer: BD

B&D because real time transfer will not have any latency  
upvoted 1 times

chiiiweiii 1 year, 6 months ago

Selected Answer: BD

Correct Answer.  
upvoted 1 times

OrangeSG 1 year, 9 months ago

Selected Answer: BD

Real time processing deals with streams of data that are captured in real-time and processed with minimal latency to generate real-time (or near-real-time) reports or automated responses.

For example, a real-time traffic monitoring solution might use sensor data to detect high traffic volumes. This data could be used to dynamically update a map to show congestion, or automatically initiate high-occupancy lanes or other traffic management systems.

Reference

Real-time processing

<https://learn.microsoft.com/en-us/azure/architecture/data-guide/big-data/real-time-processing>

upvoted 3 times

Samuel77 1 year, 9 months ago

B & D is the right answer  
upvoted 1 times

JA2018 1 year, 11 months ago

Selected Answer: BD

Data is processed in real-time as and when it is received. Low latency is a real reason why organizations opt for streaming processing.  
upvoted 4 times

DRAG DROP -

Match the Azure Data Factory components to the appropriate descriptions.

To answer, drag the appropriate component from the column on the left to its description on the right. Each component may be used once, more than once, or not at all.

NOTE: Each correct match is worth one point.

Select and Place:

Components	Answer Area
Dataset	Component A representation of data structures within data stores
Linked service	Component The information used to connect to external resources
Mapping data flow	Component A logical grouping of activities that performs a unit of work and can be scheduled
Pipeline	

Suggested Answer:

Components	Answer Area
Dataset	Dataset A representation of data structures within data stores
Linked service	Linked service The information used to connect to external resources
Mapping data flow	
Pipeline	Pipeline A logical grouping of activities that performs a unit of work and can be scheduled

Box 1: Dataset -

Datasets must be created from paths in Azure datastores or public web URLs, for the data to be accessible by Azure Machine Learning.

Box 2: Linked service -

Linked services are much like connection strings, which define the connection information needed for Data Factory to connect to external resources.

Box 3: Pipeline -

A pipeline is a logical grouping of activities that together perform a task.

Reference:

<https://k21academy.com/microsoft-azure/dp-100/datastores-and-datasets-in-azure/> <https://docs.microsoft.com/en-us/azure/data-factory/concepts-linked-services> <https://docs.microsoft.com/en-us/azure/data-factory/concepts-pipelines-activities>

 **Iroshan** Highly Voted 3 years, 5 months ago  
Correct answer, although a trick question.  
upvoted 70 times

 **Armandoo** Highly Voted 3 years, 3 months ago  
Dataset  
Linked Service  
Pipeline  
upvoted 37 times

 **AGTraining** Most Recent 6 months, 2 weeks ago

Dataset  
Linked Service  
Pipeline  
upvoted 1 times

🗄️ 👤 **TKC\_2023** 1 year ago

Dataset  
Linked Service  
Pipeline  
upvoted 1 times

🗄️ 👤 **sanmighty** 1 year, 5 months ago

Correct answer  
upvoted 2 times

🗄️ 👤 **chiiiweiii** 1 year, 6 months ago

Correct Answer.  
upvoted 1 times

🗄️ 👤 **Eltooth** 2 years, 2 months ago

Dataset  
Linked Service  
Pipeline  
upvoted 6 times

🗄️ 👤 **Kiki\_25** 2 years, 3 months ago

Correct  
upvoted 4 times

🗄️ 👤 **GP2000** 2 years, 4 months ago

Dataset, Linked Service, Pipeline  
upvoted 4 times

🗄️ 👤 **bfan95** 2 years, 4 months ago

Dataset, Linked Service, Pipeline  
upvoted 2 times

🗄️ 👤 **SiDoCiOuS** 2 years, 5 months ago

Appeared on the exam on 03/22/2022  
Answers are correct!  
upvoted 6 times

🗄️ 👤 **RajeshAzure** 2 years, 6 months ago

correct answer  
upvoted 1 times

🗄️ 👤 **Abhilash\_KK** 2 years, 8 months ago

correct answers :)  
upvoted 1 times

🗄️ 👤 **Meyti** 2 years, 9 months ago

Correct.  
Azure Data Factory and Azure Synapse Analytics:

A Data Factory or Synapse Workspace can have one or more pipeline.

Pipeline:

A pipeline is a logical grouping of activities that together perform a task and can be scheduled. The activities in a pipeline define actions to perform on your data.

a Dataset is a named view of data that simply points or references the data you want to use in your activities as inputs and outputs. Datasets identify data within different data stores, such as tables, files, folders, and documents.

Before you create a dataset, you must create a linked service to link your data store to the service. Linked services are much like connection strings, which define the connection information needed for the service to connect to external resources.

Think of it this way; the dataset represents the structure of the data within the linked data stores, and the linked service defines the connection

to the data source.

<https://docs.microsoft.com/en-us/azure/data-factory/concepts-linked-services#overview>

upvoted 13 times

🗨️ 👤 **Mjwc** 2 years, 11 months ago

Correct.

Sources:

<https://docs.microsoft.com/en-us/azure/data-factory/introduction>

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

upvoted 1 times

🗨️ 👤 **Afamdi** 3 years ago

Tricky one here

upvoted 1 times

🗨️ 👤 **Crazy\_Angel** 3 years, 2 months ago

Dataset: A data set (or dataset) is a collection of data. In the case of tabular data, a data set corresponds to one or more database tables, where every column of a table represents a particular variable, and each row corresponds to a given record of the data set in question.

Linked service: the information used to connect to external resources

Pipeline: A logical grouping of activities that performs a unit of work and can be scheduled.

upvoted 5 times

DRAG DROP -

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.

Select and Place:

### Workload Types

Batch

Streaming

### Answer Area

Workload type

Data for a product catalog will be loaded every 12 hours to a data warehouse.

Workload type

Thousands of data sets per second for online purchases will be loaded into a data warehouse in real time

Workload type

Updates to inventory data will be loaded to a data warehouse every 1 million transactions.

### Suggested Answer:

#### Workload Types

Batch

Streaming

#### Answer Area

Batch

Data for a product catalog will be loaded every 12 hours to a data warehouse.

Streaming

Thousands of data sets per second for online purchases will be loaded into a data warehouse in real time

Batch

Updates to inventory data will be loaded to a data warehouse every 1 million transactions.

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 -

Reference:

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

 **AzureDP900** Highly Voted 2 years, 11 months ago

Batch, Streaming, Batch  
upvoted 62 times

 **D\_Mo\_ZY** 2 years, 6 months ago

Was in the exam (18/03/2022)  
upvoted 12 times

 **Eltooth** Highly Voted 2 years, 2 months ago

Batch  
Streaming  
Batch  
upvoted 11 times

 **profesorklaus** Most Recent 1 week, 3 days ago

Corretc

upvoted 1 times

🗨️ 👤 **AGTraining** 6 months, 2 weeks ago

Batch, Streaming, Batch

upvoted 1 times

🗨️ 👤 **trashbox** 10 months, 2 weeks ago

Batch, Streaming, Batch.

Appeared on 2023-11-05.

upvoted 1 times

🗨️ 👤 **VNR25** 1 year, 2 months ago

This question appeared in the 07/07/23 exam and the answer is correct.

upvoted 2 times

🗨️ 👤 **PhyoPluto** 1 year, 6 months ago

the exam on 07/03/2023

upvoted 5 times

🗨️ 👤 **chiiiweiii** 1 year, 6 months ago

Correct Answer.

upvoted 3 times

🗨️ 👤 **FoxDumpz** 2 years, 1 month ago

Correct Answer

upvoted 6 times

🗨️ 👤 **heenak55** 2 years, 4 months ago

In exam 21 April 2022

upvoted 6 times

🗨️ 👤 **Ganaps** 2 years, 5 months ago

correct answer.

upvoted 2 times

🗨️ 👤 **samxhander** 2 years, 9 months ago

correct answer.

upvoted 4 times

🗨️ 👤 **Sam\_Zee** 2 years, 9 months ago

Batch, Streaming, Batch is the correct answer

upvoted 3 times

🗨️ 👤 **MrSunday** 2 years, 10 months ago

batch = conditional

stream = streaming

upvoted 3 times

🗨️ 👤 **covillmail** 2 years, 10 months ago

answer is ok

upvoted 2 times

DRAG DROP -

Your company plans to load data from a customer relationship management (CRM) system to a data warehouse by using an extract, load, and transform (ELT) process.

Where does data processing occur for each stage of the ELT process? To answer, drag the appropriate locations to the correct stages. Each location 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:

**Locations**

**Answer Area**

A standalone data analysis tool

Extract:

The CRM system

Load:

The data warehouse

Transform:

Suggested Answer:

**Locations**

**Answer Area**

Extract:

The CRM system

Load:

The data warehouse

Transform:

A standalone data analysis tool

Box 1: The CRM system -

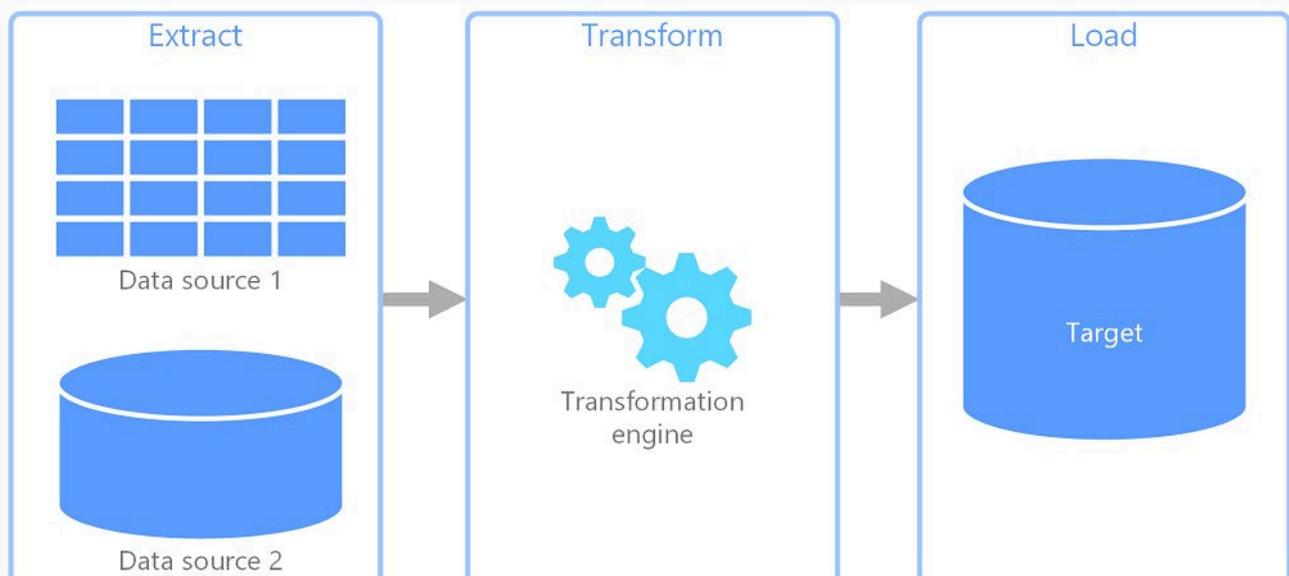
Data is extracted from the CRM system.

Box 2: The data warehouse -

Data is loaded to the data warehouse.

Box 3: A standalone data analysis tool

The data transformation that takes place usually involves various operations, such as filtering, sorting, aggregating, joining data, cleaning data, deduplicating, and validating data.



Reference:

<https://docs.microsoft.com/en-us/azure/architecture/data-guide/relational-data/etl>

🗨️ **jamorey** Highly Voted 2 years, 11 months ago

I think is CRM - DW -DW ¿is this de correct answer?

upvoted 165 times

🗨️ **bpxgeek** 2 years, 5 months ago

Yes you are right it should be CRM-DW-DW

upvoted 10 times

🗨️ **durak** 2 years, 10 months ago

Absolutely

upvoted 10 times

🗨️ **AngelJP** 2 years, 7 months ago

Yes, sure

upvoted 6 times

🗨️ **Tonnie85** 1 year, 10 months ago

Indeed...answer states for ETL, yet question specifies ELT

upvoted 4 times

🗨️ **Ezther** Highly Voted 2 years, 11 months ago

CRM-DW-DW, same question as 13

upvoted 50 times

🗨️ **wth1234** 2 years, 9 months ago

this question refers to ELT the 13th question refers to ETL

upvoted 4 times

🗨️ **wth1234** 2 years, 9 months ago

correction: you are right it does not refer to ETL oops same question

upvoted 4 times

🗨️ **rajat7878** 2 years, 6 months ago

Agreed

upvoted 1 times

🗨️ **cloudbaron** Most Recent 3 months, 1 week ago

CRM - DW- DW.

" The key point with ELT is that the data store used to perform the transformation is the same data store where the data is ultimately consumed."

Ref: <https://learn.microsoft.com/en-us/azure/architecture/data-guide/relational-data/etl>

upvoted 1 times

🗨️ **AGTraining** 6 months, 2 weeks ago

crm dw standalone

upvoted 1 times

🗨️ **Edgy\_San** 7 months, 1 week ago

The answer is correct. But the image is of an ETL process and the process in question and answer is about an an ELT. We answered this question correctly before and it had a link to read more on. Which I can't see it now, since I am here typing this.

upvoted 2 times

🗨️ **iza\_li** 7 months, 3 weeks ago

For ELT, the correct answer is CRM-DWH-DWH

<https://learn.microsoft.com/en-us/azure/architecture/data-guide/relational-data/etl>

upvoted 1 times

🗨️ **Ruesam** 10 months, 2 weeks ago

The answer here is for ETL. Correct answer for ELT should be source, DW, then DW since DW has processing power that can fit both the T and L processes within it

upvoted 1 times

🗨️ **lol632** 11 months ago

It says ELT, where your loading and transformation both happens in the Data Warehouse.

upvoted 1 times

🗨️ **MannyJZ** 1 year, 1 month ago

On exam today 08/08/23

upvoted 2 times

🗨️ **fafulka** 1 year, 3 months ago

in exam 16/06/2023

upvoted 2 times

🗨️ **Dudures** 1 year, 5 months ago

Exam 04/15/2023

upvoted 4 times

🗨️ **chiiiweiii** 1 year, 6 months ago

Correct Answer.

upvoted 1 times

🗨️ **Mohitsain** 1 year, 7 months ago

1.CRM

2.DW

3.DW

upvoted 2 times

🗨️ **Stuurwiel1** 1 year, 7 months ago

ELT> CRM> DW>DW

ETL> CRM> DW> DAT

<https://learn.microsoft.com/en-us/azure/architecture/data-guide/relational-data/etl>

upvoted 2 times

🗨️ **cpalmier** 1 year, 9 months ago

Pienso que es correcta la respuesta, si considero a Power BI. Power BI es una herramienta de análisis la cual puede hacer transformaciones.

upvoted 1 times

🗨️ **Dusica** 1 year, 10 months ago

CRM, DW, DW

upvoted 2 times

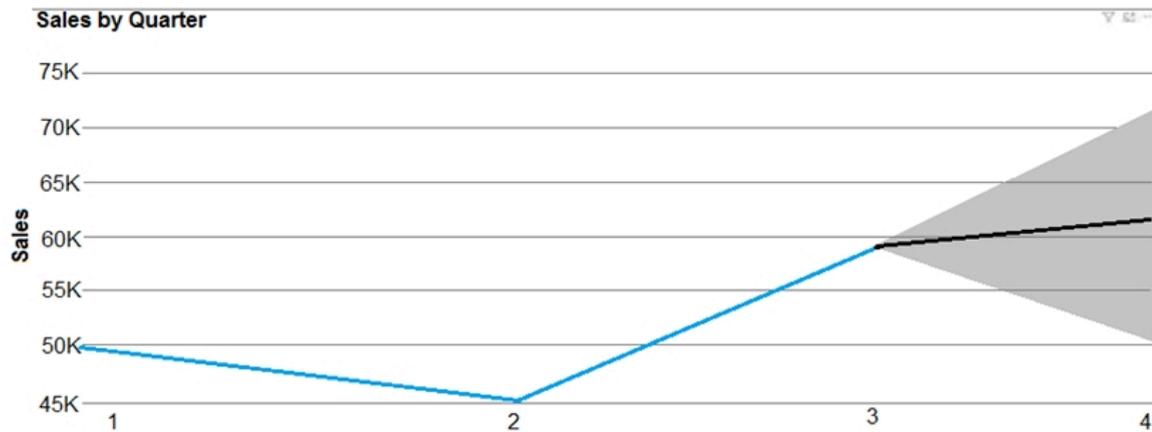
🗨️ **tellwade** 1 year, 12 months ago

The tool here is different to question 13. in an ELT you can extract from the CRM, load to the DW and Transform in a tool like power BI? very confusing not sure what's correct

upvoted 1 times

Your company recently reported sales from the third quarter.

You have the chart shown in the following exhibit.



Which type of analysis is shown in the fourth quarter?

- A. predictive
- B. prescriptive
- C. descriptive
- D. diagnostic

**Suggested Answer: A**

Predictive, to answer the question: What will happen?



Reference:

<https://demand-planning.com/2020/01/20/the-differences-between-descriptive-diagnostic-predictive-cognitive-analytics/>

<https://azure.microsoft.com/en-us/blog/answering-whats-happening-whys-happening-and-what-will-happen-with-iot-analytics/>

Community vote distribution



**PanP** Highly Voted 2 years, 10 months ago

If just Q1-Q3 -> Descriptive

Include Q3-Q4 -> become Predictive

upvoted 47 times

**pukso metron** Highly Voted 2 years, 11 months ago

Predictive seems appropriate

upvoted 29 times

**juanl91** Most Recent 4 months, 1 week ago

**Selected Answer: A**

A is correct

upvoted 1 times

AGTraining 6 months, 2 weeks ago

Selected Answer: A

option A

upvoted 1 times

riwahof 11 months, 3 weeks ago

Selected Answer: A

Correct

upvoted 1 times

kjszala 11 months, 3 weeks ago

Selected Answer: A

correct

upvoted 1 times

vicks60 1 year, 5 months ago

Selected Answer: A

A is correct ans.

upvoted 2 times

chiiiweiii 1 year, 6 months ago

Selected Answer: A

Correct Answer.

upvoted 2 times

Joanna0 1 year, 7 months ago

Selected Answer: A

Predictive, because Q3-Q4 is in the future and they will try to predict that sales may be in a certain scope.

upvoted 1 times

Glubbs 1 year, 10 months ago

Selected Answer: A

A is the correct answer (Q3-Q4)

upvoted 3 times

JA2018 1 year, 11 months ago

Selected Answer: A

Predict what might happen in Q4

upvoted 3 times

raydel92 1 year, 11 months ago

Selected Answer: A

Predicts what will happen

upvoted 3 times

xRiot007 2 years, 2 months ago

A - predictive analysis. Your last report was for Q3 and based on Q1-Q3 data, you estimate what Q4 might be.

upvoted 2 times

Eltooth 2 years, 2 months ago

Selected Answer: A

A is correct answer.

upvoted 1 times

EXODUS33 2 years, 4 months ago

Selected Answer: A

predictive as its for the future

upvoted 2 times

RichMelbourne 2 years, 5 months ago

Shown on test 15/04/2022

upvoted 5 times

SiDoCiOuS 2 years, 5 months ago

Selected Answer: A

Appeared on the exam on 03/22/2022

Answer is correct!

upvoted 4 times

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

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

**Suggested Answer:** D

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 -

Reference:

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

Community vote distribution

D (100%)

 **Roy\_zuniga** Highly Voted 2 years, 8 months ago

There's 3 type of Relational data SQL:

DDL (Data Definition Language) CREATE, DROP, ALTER, RENAME, COMMENT AND TRUNCATE.

DML (Data Manipulación Language) SELECT, INSERT INTO, DELETE, UPDATE.

DCL (Data Control Language) REVOKE & GRANT.

TCL (Transaccional Control Language) COMMIT & ROLLBACK

(Comment from a previous question.)

upvoted 38 times

 **Nurialzard** 1 year, 7 months ago

Select is DQL

upvoted 1 times

 **Loofy** Highly Voted 2 years, 11 months ago

Answer is D' Update

upvoted 24 times

 **yusufuthman57** 2 years, 10 months ago

D is correct answer

UPDATE

upvoted 6 times

 **juanl91** Most Recent 4 months, 1 week ago

**Selected Answer: D**

D is the correct

upvoted 1 times

 **AGTraining** 6 months, 2 weeks ago

option D

upvoted 1 times

🗨️ **kjszala** 11 months, 3 weeks ago

**Selected Answer: D**

correct

upvoted 1 times

🗨️ **akhil5432** 1 year, 1 month ago

Option D

upvoted 1 times

🗨️ **VNR25** 1 year, 2 months ago

This question appeared in the 07/07/23 exam and the answer is correct.

upvoted 1 times

🗨️ **fafulka** 1 year, 3 months ago

in exam 16/06/2023

upvoted 1 times

🗨️ **chiiiweiii** 1 year, 6 months ago

**Selected Answer: D**

Correct Answer.

upvoted 2 times

🗨️ **Metarineo** 1 year, 10 months ago

@Roy\_zuniga:

DDL (Data Definition Language) CREATE, DROP, ALTER, RENAME, COMMENT AND TRUNCATE.

DML (Data Manipulation Language) SELECT, INSERT INTO, DELETE, UPDATE.

I have a conflict with the "TRUNCATE" as being DDL. I would rather put it into DML. Does MS have a clear information about this?

upvoted 2 times

🗨️ **JA2018** 1 year, 11 months ago

**Selected Answer: D**

You require DML to modify the information stored in the database.

upvoted 1 times

🗨️ **raydel92** 1 year, 11 months ago

**Selected Answer: D**

D, update

upvoted 1 times

🗨️ **john1979** 1 year, 11 months ago

**Selected Answer: D**

D as by this you update the transaction (Data Manipulation Language) so updating is manipulation

upvoted 1 times

🗨️ **JftCoCo** 2 years ago

**Selected Answer: D**

Create -> Insert, Not Create

upvoted 2 times

🗨️ **T4321** 2 years, 1 month ago

**Selected Answer: D**

Answer is 'D'

upvoted 1 times

🗨️ **Drchattss** 2 years, 1 month ago

01/08.....EXAM

upvoted 4 times

🗨️ **Eltooth** 2 years, 2 months ago

**Selected Answer: D**

D is correct answer.

upvoted 2 times



HOTSPOT -

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

Hot Area:

## Answer Area

Creating closed caption text for audio files is an example of  analytics.

<input type="text"/>	▼
cognitive	
descriptive	
predictive	
prescriptive	

Suggested Answer:

## Answer Area

Creating closed caption text for audio files is an example of  analytics.

<input type="text"/>	▼
cognitive	
descriptive	
predictive	
prescriptive	

Descriptive, to answer the question: What's happening?

			
<b>Descriptive</b> <i>What's happening?</i>	<b>Diagnostic</b> <i>Why's happening?</i>	<b>Predictive</b> <i>What will happen?</i>	<b>Prescriptive</b> <i>What actions should we take?</i>

Note: Azure Media Indexer enables you to make content of your media files searchable and to generate a full-text transcript for closed captioning and keywords.

You can process one media file or multiple media files in a batch.

Reference:

<https://demand-planning.com/2020/01/20/the-differences-between-descriptive-diagnostic-predictive-cognitive-analytics/>

<https://azure.microsoft.com/en-us/blog/answering-whats-happening-whys-happening-and-what-will-happen-with-iot-analytics/>

<https://docs.microsoft.com/en-us/azure/media-services/previous/media-services-index-content>

 **Fabsworth** Highly Voted 2 years, 4 months ago

I'm almost certain this is Cognitive.

upvoted 212 times

 **Lscrinio** 10 months ago

Yes, It's Cognitive.

upvoted 5 times

 **ha1p** 1 year, 5 months ago

agree w/ you

upvoted 3 times

 **ColoradoChick** 1 year, 2 months ago

Yes - it's Cognitive. Closed captioning is speech-to-text, which is in Azure Cognitive Services Speech Service. And the tool mentioned in the explanation on the web page has been retired (according to MS Docs)

upvoted 13 times

  **pads98** 2 years, 4 months ago

Answer is definitely Cognitive.

<https://azure.microsoft.com/en-us/services/cognitive-services/speech-to-text/#overview>

upvoted 43 times

  **Mannie\_001**  1 year, 10 months ago

COGNITIVE

upvoted 20 times

  **profesorklaus**  1 week, 3 days ago

Cognitive is a right answer.

upvoted 1 times

  **TKC\_2023** 1 month, 1 week ago

Cognitive

upvoted 1 times

  **09c92f8** 1 month, 1 week ago

100 % cognitive

upvoted 1 times

  **okayhey** 6 months, 1 week ago

all of us in the comments are smarter than the person uploading there answers xD bit appreciate them uploading it tho

upvoted 2 times

  **AGTraining** 6 months, 2 weeks ago

Cognitive

upvoted 1 times

  **jxs221** 7 months ago

Creating closed caption text for audio files is an example of cognitive analytics. This involves the use of AI and machine learning technologies to simulate human processes such as understanding language and recognizing speech.

upvoted 3 times

  **Edgy\_San** 7 months, 1 week ago

I chose cognitive. Not sure how descriptive would relate to the question and it not asking what happened.

upvoted 1 times

  **BXtheCoder** 8 months ago

Isn't it Cognitive?

upvoted 1 times

  **ff09375** 8 months, 1 week ago

FWIW, from ChatGPT:

Creating closed caption text for audio files is an example of cognitive analytics. Cognitive analytics involves using machine learning and natural language processing (NLP) technologies to automate the processing of large amounts of data, including audio and text data, to derive insights and improve decision-making. In this case, it's used to transcribe spoken words into written text. This type of analytics is particularly useful in applications like speech recognition, sentiment analysis, and language translation.

upvoted 3 times

  **Rudramina** 9 months, 3 weeks ago

Descriptive

upvoted 2 times

  **riwahof** 11 months, 3 weeks ago

Bing (with precision mode) gives the following:

Creating closed caption text for audio files is an example of cognitive analytics. Cognitive analytics involves applying intelligent algorithms to interpret unstructured data like text, speech, or images. In this case, it's interpreting audio data and converting it into text.

upvoted 4 times

  **Law3** 11 months, 4 weeks ago

I think we are not talking about the cognitive transcription process, but what happen next, as you create the caption for accessibility and indexing. You don't need to understand the audio file, you need to make it searchable. Can the user directly fill the closed caption by himself, without using any speech-to-text ?

upvoted 1 times

  **Prince9598** 1 year ago

Cognitive

upvoted 1 times

  **hcq31818** 1 year ago

It is cognitive

upvoted 1 times

  **akhil5432** 1 year, 1 month ago

descriptive

upvoted 2 times

HOTSPOT -

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

Hot Area:

### Answer Area

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

<input type="text"/>	▼
cognitive	
descriptive	
predictive	
prescriptive	

### Answer Area

Suggested Answer: A visualization that illustrates a university's current student enrollment per department is an example of  analytics.

<input type="text"/>	▼
cognitive	
descriptive	
predictive	
prescriptive	

Generally speaking, data analytics comes in four types:

1. Descriptive, to answer the question: What's happening?
2. Diagnostic, to answer the question: Why's happening?
3. Predictive, to answer the question: What will happen?
4. Prescriptive, to answer the question: What actions should we take?



Reference:

<https://azure.microsoft.com/en-us/blog/answering-whats-happening-whys-happening-and-what-will-happen-with-iot-analytics/>

🗨️ **riglva** Highly Voted 2 years, 10 months ago

Yes, it's descriptive  
upvoted 48 times

🗨️ **Srismiles** Highly Voted 2 years ago

Yes. Descriptive  
upvoted 11 times

🗨️ **AGTraining** Most Recent 6 months, 2 weeks ago

Descriptive  
upvoted 1 times

🗨️ **james2033** 10 months, 3 weeks ago

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

upvoted 1 times

🗨️ **chiiiweiii** 1 year, 6 months ago

Correct Answer.

upvoted 2 times

🗨️ **iwikneerg** 2 years, 1 month ago

descriptive because it is showing what is

upvoted 7 times

🗨️ **Eltooth** 2 years, 2 months ago

Descriptive

upvoted 4 times

🗨️ **GP2000** 2 years, 4 months ago

Descriptive

upvoted 5 times

🗨️ **RichMelbourne** 2 years, 5 months ago

Shown on test 15/04/2022

upvoted 5 times

🗨️ **meldr** 2 years, 5 months ago

Correct, descriptive

upvoted 3 times

🗨️ **Ganaps** 2 years, 5 months ago

CORRECT

upvoted 1 times

🗨️ **SiDoCiOuS** 2 years, 5 months ago

Appeared on the exam on 03/22/2022

Answer is correct!

upvoted 3 times

🗨️ **Roy\_zuniga** 2 years, 8 months ago

Correct.

1. Descriptive, to answer the question: What's happening?
2. Diagnostic, to answer the question: Why's happening?
3. Predictive, to answer the question: What will happen?
4. Prescriptive, to answer the question: What actions should we take?

upvoted 3 times

🗨️ **Linceyl** 2 years, 8 months ago

Correct

upvoted 1 times

🗨️ **jqx60863** 2 years, 9 months ago

The same question as before, and few others

upvoted 1 times

🗨️ **Abh123** 2 years, 10 months ago

CORRECT

upvoted 3 times

You need to create an Azure resource to store data in Azure Table storage.

Which command should you run?

- A. az storage share create
- B. az storage account create
- C. az cosmosdb create
- D. az storage container create

**Suggested Answer:** D

Reference:

<https://docs.microsoft.com/en-us/cli/azure/storage/container?view=azure-cli-latest>

Community vote distribution



**PanP** Highly Voted 2 years, 10 months ago

Ans is B for me. Due to no az storage table create in the selection. The best answer is B create storage account (first)  
upvoted 82 times

**Kanatius** 1 year, 3 months ago

I don't even think that a container is considered as an "Azure Resource" (the Azure Resources come right after the resource group), so it should be storage account (an Azure Resource)  
upvoted 3 times

**amraneze** 2 years, 7 months ago

Normally, they should write table but if it's not mentioned then account because you can't have a table in azure table without an azure storage account  
upvoted 10 times

**walkwolf3** Highly Voted 2 years, 10 months ago

It could be "az storage table create"

<https://docs.microsoft.com/en-us/cli/azure/storage/table?view=azure-cli-latest>

upvoted 16 times

**profesorklaus** Most Recent 1 week, 3 days ago

First you need to create a storage so B is a correct answer  
upvoted 1 times

**Ruesam** 4 months ago

From meta AI:

To create an Azure resource for storing data in Azure Table storage, you can use the Azure CLI command:

```
`az storage account create --name <storage-account-name> --resource-group <resource-group-name> --location <location> --sku Standard_LRS`
```

Replace ``<storage-account-name>`` with the desired name for your storage account, ``<resource-group-name>`` with the name of the resource group, and ``<location>`` with the desired location (such as "West US" or "East US").

Once the storage account is created, you can create a table using the command:

```
`az storage table create --name <table-name> --account-name <storage-account-name>`
```

Replace ``<table-name>`` with the desired name for your table, and ``<storage-account-name>`` with the name of the storage account you created.  
upvoted 1 times

**AGTraining** 6 months, 2 weeks ago

az storage account create

upvoted 3 times

🗨️ 👤 **qasmi** 7 months, 2 weeks ago

**Selected Answer: B**

To create resource first account needs to be created.

upvoted 2 times

🗨️ 👤 **BrennoVisk** 8 months, 2 weeks ago

**Selected Answer: B**

Reply Chat gpt

I apologize for the mistake in my previous answer. If option D is associated with creating storage containers in the context of Azure Table Storage, this would be an incorrect answer. Azure tables are part of the Azure Table Storage service, which does not use containers like Azure Blob Storage.

The correct answer to storing data in Azure Table Storage would be to create an Azure Storage Account. Therefore, the correct option would be the letter B:

upvoted 4 times

🗨️ 👤 **masadamsahid** 10 months, 3 weeks ago

Azure Table Storage doesn't need a container. AFAIK, container is created for a BLOB storage. Since Azure Table Storage isn't a BLOB storage, I don't think choice D is the right answer.

And since there is no "az storage table create" choice, the most appropriate choice should be choice B. It's pretty reasonable since we need to create a subscription account first.

upvoted 2 times

🗨️ 👤 **james2033** 10 months, 3 weeks ago

**Selected Answer: B**

az storage account create

<https://learn.microsoft.com/en-us/cli/azure/storage/account?view=azure-cli-latest#az-storage-account-create>

upvoted 2 times

🗨️ 👤 **nserge** 10 months, 4 weeks ago

B. Create an account

<https://learn.microsoft.com/en-us/azure/storage/tables/table-storage-how-to-use-powershell>

upvoted 2 times

🗨️ 👤 **mrharveypotter** 11 months, 1 week ago

**Selected Answer: D**

Totally agree with Trulysme, option "D".

upvoted 2 times

🗨️ 👤 **Trulysme** 11 months, 2 weeks ago

Read the question carefully. It is intended to be confusing!

It does ask to create a Azure Ressource within an Azure Table Storage. Which indicates that the account is allready created. In that case you would create a ressource which would be a CONTAINER. Meaning answer "D" is indeed the correct answer.

upvoted 4 times

🗨️ 👤 **masadamsahid** 10 months, 3 weeks ago

AFAIK, container is only for a BLOB storage. Never heard a container could be created for a NoSQL resource such as Azure Table Storage.

upvoted 1 times

🗨️ 👤 **Ifparauj0** 1 year ago

According to Chat GPT:

To create an Azure resource to store data in Azure Table storage, you should run the following command:

B. az storage account create

Creating an Azure Storage Account is the initial step in setting up Azure Table storage. Once you have created the storage account, you can then work with tables and containers within that account to store and manage your data. Options A, C, and D do not directly create a storage account for Azure Table storage.

upvoted 1 times

🗨️ 👤 **AZFabio** 1 year ago

**Selected Answer: B**

Create a storage account, then a table within it

upvoted 1 times

🗨️ 👤 **mdkshah** 1 year ago

**Selected Answer: B**

To store data in Azure Table storage, you need to create an Azure Storage account first. You can create an Azure Storage account using the command `az storage account create`. This command creates a new storage account in Azure with the specified name, resource group, location, and other configurations such as the performance tier, replication type, and access control settings

upvoted 1 times

🗨️ 👤 **akhil5432** 1 year, 1 month ago

option D

upvoted 1 times

🗨️ 👤 **sharmab5** 1 year, 1 month ago

B is correct

upvoted 1 times

HOTSPOT -

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

Hot Area:

### 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.

Suggested Answer:

### 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.

Azure Cosmos DB SQL API accounts provide support for querying items using the Structured Query Language (SQL) syntax.

Reference:

<https://docs.microsoft.com/en-us/azure/cosmos-db/choose-api>

 **Abhijinendra** Highly Voted 2 years, 10 months ago

Core(SQL)

upvoted 57 times

 **Koph** 2 years, 9 months ago

Agree, answer is correct

upvoted 6 times

 **mamahani** Highly Voted 1 year, 5 months ago

I took the exam 12.04.2023 and got 1000, 100% score; this question appeared, correct answer:

Core (SQL)

upvoted 10 times

 **AGTraining** Most Recent 6 months, 2 weeks ago

Core(SQL)

upvoted 1 times

 **chiiweiii** 1 year, 6 months ago

Correct Answer

upvoted 2 times

 **KADANGI** 1 year, 10 months ago

just passed with 833 but I get first option as NoSQL

upvoted 6 times

 **Kanatius** 1 year, 4 months ago

By default cosmosDB stores data as json and you can retrieve data using select/where, so the first option is the correct one

upvoted 2 times

🗨️ 👤 **Css6000** 1 year, 5 months ago

So nosql would be the correct answer right?

upvoted 1 times

🗨️ 👤 **s3s** 1 year, 8 months ago

Same with me on 04/01/2022

upvoted 2 times

🗨️ 👤 **Kermedi2018** 1 year, 11 months ago

This is Correct Answer it came on my exam (03/10/2022) , I just passed my Exam Thank You Exam Topic group i got 833

upvoted 3 times

🗨️ 👤 **blackskies** 1 year, 11 months ago

was on exam today 09/30/22

upvoted 5 times

🗨️ 👤 **Eltooth** 2 years, 2 months ago

Core is correct answer.

upvoted 1 times

🗨️ 👤 **dk87** 2 years, 4 months ago

Exam: 05/18/2022

upvoted 4 times

🗨️ 👤 **RichMelbourne** 2 years, 5 months ago

Shown on test 15/04/2022

upvoted 9 times

🗨️ 👤 **waxun** 2 years, 5 months ago

@RajeshAzure Because it uses the select statement ( sql)

upvoted 2 times

🗨️ 👤 **SiDoCiOus** 2 years, 5 months ago

Appeared on the exam on 03/22/2022

Answer is correct!

upvoted 4 times

🗨️ 👤 **RajeshAzure** 2 years, 6 months ago

Mongo DB API for the the Document storage, should not it be Mongo DB or what am i missing ?

upvoted 1 times

🗨️ 👤 **BabyBee** 2 years, 5 months ago

No, because they are asking to use a SELECT.

Example of a query with Mongo: `db.families.find( { id: "WakefieldFamily" }, { children: true } )`

upvoted 1 times

🗨️ 👤 **RichieB** 2 years, 7 months ago

Core (SQL)

upvoted 1 times

🗨️ 👤 **MohammadKhubeb** 2 years, 8 months ago

MongoDB

upvoted 4 times

🗨️ 👤 **erhard** 1 year, 10 months ago

MongoDB is not MariaDB.

upvoted 1 times

🗨️ 👤 **Artsem** 2 years, 8 months ago

Right one

upvoted 1 times

You need to modify a view in a relational database by adding a new column.  
Which statement should you use?

- A. MERGE
- B. ALTER
- C. INSERT
- D. UPDATE

**Suggested Answer:** B

Community vote distribution

B (100%)

-  **theLazyGuy** Highly Voted 2 years, 10 months ago  
to add a column in a view Alter (<https://docs.microsoft.com/en-us/sql/relational-databases/tables/add-columns-to-a-table-database-engine?view=sql-server-ver15>)  
Insert adds one or more rows to a table (<https://docs.microsoft.com/en-us/sql/t-sql/statements/insert-transact-sql?view=sql-server-ver15>)  
upvoted 27 times
-  **AdminFromHell** Highly Voted 2 years, 9 months ago  
Selected Answer: B  
"Modify a View" = Alter View ....  
upvoted 25 times
-  **AGTraining** Most Recent 6 months, 2 weeks ago  
option B  
upvoted 1 times
-  **fe6d44f** 8 months, 1 week ago  
Passed exam on 9/1/2024, got 984/1000.  
My answer was ALTER  
upvoted 3 times
-  **kjszala** 11 months, 3 weeks ago  
Selected Answer: B  
Alter (can be also Create or Alter)  
upvoted 1 times
-  **MannyJZ** 1 year, 1 month ago  
On exam today 08/08/23  
upvoted 3 times
-  **elbert92** 1 year, 1 month ago  
Selected Answer: B  
Alter to add a column  
upvoted 1 times
-  **Atif1977** 1 year, 2 months ago  
Selected Answer: B  
Alter to add a column  
upvoted 1 times
-  **you5uf** 1 year, 3 months ago  
in exam it will come as NoSQL, inplace for CoreSQL, you should choose A)  
  
<https://learn.microsoft.com/en-us/azure/cosmos-db/choose-api#coresql-api>  
upvoted 1 times
-  **fifulka** 1 year, 3 months ago

in exam 16/06/2023

upvoted 1 times

🗨️ **chiiiweiii** 1 year, 6 months ago

**Selected Answer: B**

Correct Answer

upvoted 2 times

🗨️ **edurakhan** 1 year, 10 months ago

In exam today 11/15/22

upvoted 2 times

🗨️ **john1979** 1 year, 11 months ago

**Selected Answer: B**

any change in column, needs to alter the table

upvoted 4 times

🗨️ **ria97** 2 years, 1 month ago

Was in exam 11/08.

upvoted 3 times

🗨️ **Eltooth** 2 years, 2 months ago

**Selected Answer: B**

B is correct answer.

upvoted 3 times

🗨️ **muks1003** 2 years, 3 months ago

**Selected Answer: B**

Correct

upvoted 1 times

🗨️ **EXODUS33** 2 years, 4 months ago

**Selected Answer: B**

Alter is the ans

upvoted 2 times

HOTSPOT -

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE: Each correct selection is worth one point.

Hot Area:

### Answer Area

Statements	Yes	No
Extract, transform, and load (ETL) can reduce the transfer of sensitive data to destination systems.	<input type="radio"/>	<input type="radio"/>
Extract, load, and transform (ELT) transforms data by using a compute resource independent of the source system and destination system.	<input type="radio"/>	<input type="radio"/>
Extract, load, and transform (ELT) minimizes the time it takes to copy large volumes of data to destination systems.	<input type="radio"/>	<input type="radio"/>

### Answer Area

Statements	Yes	No
Extract, transform, and load (ETL) can reduce the transfer of sensitive data to destination systems.	<input type="radio"/>	<input checked="" type="radio"/>
Extract, load, and transform (ELT) transforms data by using a compute resource independent of the source system and destination system.	<input checked="" type="radio"/>	<input type="radio"/>
Extract, load, and transform (ELT) minimizes the time it takes to copy large volumes of data to destination systems.	<input type="radio"/>	<input checked="" type="radio"/>

Suggested Answer:

Reference:

<https://docs.microsoft.com/en-us/azure/architecture/data-guide/relational-data/etl>

 **Kardec** Highly Voted 2 years, 10 months ago

If you to think logically,

Sentence one is yes - You can reduce transfer sensitive data with transformations on data source.

Sentence two is No - ELT use the compute resource from destination system.

And for last Yes. ELT minimizes time with transformations are going realized on target data destination systems.

upvoted 336 times

 **duds19** 2 years, 10 months ago

ETL phases are run in parallel to save time...the last is No

upvoted 6 times

 **datakid** 2 years, 10 months ago

The last one is yes, because ELT processes skip the data copy step present in ETL, which can be a time consuming operation for large data sets.

<https://docs.microsoft.com/en-us/azure/architecture/data-guide/relational-data/etl>

upvoted 17 times

 **aurorafang** 2 years, 2 months ago

this is correct answers

upvoted 3 times

  **rajat7878** 2 years, 6 months ago

Totally agreed. Answer is Yes, No, Yes

upvoted 12 times

  **bpxgeek** 2 years, 5 months ago

Correct

upvoted 4 times

  **theLazyGuy** Highly Voted 2 years, 10 months ago

Y/N/Y

ETL: ...data transformation that takes place usually involves various operations, such as filtering...

ELT: ...This approach skips the data copy step present in ETL, which can be a time consuming operation for large data sets...

Load and Trasformation are on the same target

(<https://docs.microsoft.com/en-us/azure/architecture/data-guide/relational-data/etl>)

upvoted 92 times

  **shahrzadkhb** Most Recent 4 weeks ago

I believe the correct answer is Yes, No, Yes.

upvoted 1 times

  **jerrybin30731** 2 months, 2 weeks ago

yes, no, yes

upvoted 1 times

  **AGTraining** 6 months, 2 weeks ago

Si No Si

upvoted 3 times

  **son\_el** 6 months, 3 weeks ago

ETL can redact and remove sensitive information before putting it into the data warehouse or cloud server. so i think no1 is yes

upvoted 1 times

  **BrennoVisk** 8 months, 2 weeks ago

<https://www.striim.com/blog/etl-vs-elt-differences/>

upvoted 1 times

  **BrennoVisk** 8 months, 2 weeks ago

<https://aws.amazon.com/pt/what-is/etl/>

upvoted 1 times

  **lol632** 11 months ago

The answer for 2 is NO. The reason being, it says source and destination, those are the keywords. In ELT, the source for transformation is inside DW itself and the destination is also inside DW. So, the compute is same. Answer is NO.

The answer for 1 - Yes, as transformation allows fileting data before you put sensitive information in a data warehouse or any destination. This could be really beneficial in for example, banking sector.

The answer for 3 - Yes, and the reason is in ETL processes, one of the bottleneck was staging server compute which took a lot of time. Whereas in ELT process, the transformation takes inside DW, which supports massive parallel computing and it removes that bottleneck of slow compute.

upvoted 6 times

  **okayhey** 6 months, 1 week ago

I agree, data doesn't have to be copied again once it's in the DW, and it's also transformed in there, so no need for separate transformation tool

1 is yes, data in etl is transformed before being loaded, so could be anonymised before loading

upvoted 1 times

  **sunny12dhy** 12 months ago

it should be NO,NO,Yes

upvoted 1 times

  **rmeng** 1 year, 1 month ago

Yes, No, Yes

upvoted 1 times

  **rmeng** 1 year, 1 month ago

Yes

No

Yes

upvoted 2 times

🗨️ 👤 **happybenny** 1 year, 1 month ago

no yes yes

upvoted 1 times

🗨️ 👤 **Ahmadov\_10** 1 year, 2 months ago

Instead of using a separate transformation engine, the processing capabilities of the target data store are used to transform data.

<https://learn.microsoft.com/en-us/azure/architecture/data-guide/relational-data/etl>

the answer : NO / NO / YES

upvoted 2 times

🗨️ 👤 **Ahmadov\_10** 1 year, 2 months ago

I think No , No , Yes :

for the first sentence , i found it's crazy to say ETL can reduce transfer sensitive : it's a not a property of ETL , you should define a solution by yourself

upvoted 3 times

🗨️ 👤 **fifulka** 1 year, 3 months ago

in exam 16/06/2023

upvoted 3 times

🗨️ 👤 **Elderbds** 1 year, 3 months ago

De acordo com o BING AÍ, a única resposta correta é: "Extrair, carregar e transformar (ELT) minimiza o tempo necessário para copiar grandes volumes de dados para os sistemas de destino".

upvoted 1 times

HOTSPOT -

You plan to deploy a PostgreSQL database to Azure.

Which hosting model corresponds to the available deployment options? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

### Answer Area

PostgreSQL on Azure VMs:

Infrastructure as a service (IaaS)
Platform as a service (PaaS)
Software as a service (SaaS)

Azure Database for PostgreSQL:

Infrastructure as a service (IaaS)
Platform as a service (PaaS)
Software as a service (SaaS)

### Answer Area

Suggested Answer:

PostgreSQL on Azure VMs:

Infrastructure as a service (IaaS)
Platform as a service (PaaS)
Software as a service (SaaS)

Azure Database for PostgreSQL:

Infrastructure as a service (IaaS)
Platform as a service (PaaS)
Software as a service (SaaS)

Reference:

<https://azure.microsoft.com/en-us/overview/what-is-saas/>

<https://azure.microsoft.com/en-us/overview/what-is-paas/>

 **Kardec** Highly Voted 2 years, 10 months ago

The correct answer is iaas and paas.

<https://docs.microsoft.com/en-us/azure/postgresql/overview-postgres-choose-server-options>

upvoted 368 times

 **jellybiscuit** 2 years ago

It sure would be great if the site updated their wrong answers. 10 months later, and it's still wrong.

upvoted 38 times

 **J\_ose** 1 year, 8 months ago

Completely agree. Please show professionalism and attention to detail

upvoted 6 times

 **gabo** 7 months ago

Exactly, and we paid a lot of money to get these answers, they need to review all the answers and update with the right ones.

upvoted 2 times

 **ColoradoChick** 1 year, 2 months ago

jellybiscuit - I agree. A LOT of the "official" answers are simply wrong. But when I look at an answer and think "WTF! That's not right!" and then look at the discussion .... well, in an odd way that's testing my readiness for the exam, too.

It'll screw the people who just memorize the questions & answers (rather than those of us who learn the material, review the discussion, etc). My advice to people is know the material - because some of the answers in the discussions are wrong, too.

upvoted 6 times

  **BobOkech** 3 months, 4 weeks ago

I agree with you.

upvoted 1 times

  **SameerL** 2 years, 2 months ago

>>> PostgreSQL on Azure VMs option falls into the industry category of IaaS per below post.....

<https://docs.microsoft.com/en-us/azure/postgresql/single-server/overview-postgres-choose-server-options>

So the correct answer is IaaS and PaaS.

upvoted 14 times

  **blueCrystal** 2 years, 7 months ago

true as you can see in the document, under Postgres on Azure VMs, everything is customer managed that means Postgres on Azure VMs is an IaaS

upvoted 12 times

  **duds19**  2 years, 10 months ago

IaaS and PaaS

upvoted 55 times

  **sdsz**  4 weeks ago

IaaS and PaaS

upvoted 1 times

  **asrasm2024** 3 months ago

I guess the answers should be IaaS and PaaS

upvoted 1 times

  **Vitos25** 5 months, 1 week ago

With Azure, your PostgreSQL workloads can run in a hosted virtual machine infrastructure as a service (IaaS) or as a hosted platform as a service (PaaS).

upvoted 1 times

  **QeUtPat** 5 months, 2 weeks ago

IaaS and PaaS are the correct answers, please update it so other people don't get confused.

upvoted 2 times

  **AGTraining** 6 months, 2 weeks ago

IaaS

PaaS

upvoted 1 times

  **BXtheCoder** 7 months, 3 weeks ago

The correct answer is IaaS and PaaS. Virtual Machines (VMs) fall under the IaaS category and therefore, the correct answer is IaaS. Azure Database for PostgreSQL is PaaS, as DBaaS (Database as a Service) falls under the PaaS category. Hope this helps!

upvoted 1 times

  **ravinperera** 9 months, 3 weeks ago

The correct answer is IaaS and PaaS, let me explain, In the case of PostgreSQL on Azure VMs, we are responsible for provisioning, managing, and maintaining both the virtual machines (the infrastructure) and the PostgreSQL software itself. It falls more under the Infrastructure as a Service (IaaS) category, as we have control over the infrastructure and software configuration, but you still manage many aspects of the environment.

Azure Database for PostgreSQL falls under the category of Database as a Service (DBaaS) or Platform as a Service (PaaS).

upvoted 2 times

  **TKri** 10 months, 2 weeks ago

Answer is IaaS and PaaS

upvoted 1 times

  **hcq31818** 1 year ago

1 --> IaaS

2--> PaaS

upvoted 1 times

🗨️ 👤 **rmeng** 1 year, 1 month ago

1. PostgreSQL on Azure VMs: In this scenario, you are deploying and managing PostgreSQL on virtual machines (VMs) provided by Azure. You have full control over the VMs and the PostgreSQL installation, making it an Infrastructure as a Service (IaaS) model. You are responsible for managing the operating system, database software, and configuration.

2. Azure Database for PostgreSQL: With Azure Database for PostgreSQL, Microsoft manages the underlying infrastructure, including VMs, storage, and backups. You only need to focus on managing the PostgreSQL databases and the data itself. This service abstracts the infrastructure layer, providing a managed database service, making it a Platform as a Service (PaaS) model.

In summary, when you use PostgreSQL on Azure VMs, it is an IaaS model, while Azure Database for PostgreSQL is a PaaS model.

upvoted 3 times

🗨️ 👤 **Myrish** 1 year, 2 months ago

IaaS and PaaS

PostgreSQL on Azure VMs option falls into the industry category of IaaS. With this service, you can run PostgreSQL Server inside a fully managed virtual machine on the Azure cloud platform. All recent versions and editions of PostgreSQL can be installed on an IaaS virtual machine.

Source: <https://learn.microsoft.com/en-us/azure/postgresql/flexible-server/overview-postgres-choose-server-options>

upvoted 1 times

🗨️ 👤 **GermanGerman** 1 year, 2 months ago

VMs--IaaS,

upvoted 1 times

🗨️ 👤 **MacHayward** 1 year, 3 months ago

PostgreSQL on Azure VMs option falls into the industry category of IaaS. With this service, you can run PostgreSQL Server inside a fully managed virtual machine on the Azure cloud platform. All recent versions and editions of PostgreSQL can be installed on an IaaS virtual machine. In the most significant difference from Azure Database for PostgreSQL, PostgreSQL on Azure VMs offers control over the database engine. However, this control comes at the cost of responsibility to manage the VMs and many database administration (DBA) tasks. These tasks include maintaining and patching database servers, database recovery, and high-availability design.

<https://learn.microsoft.com/en-us/azure/postgresql/flexible-server/overview-postgres-choose-server-options>

upvoted 1 times

🗨️ 👤 **Elderbds** 1 year, 4 months ago

De acordo com o BING IA da Microsoft: Sim, executar o PostgreSQL em uma máquina virtual (VM) do Azure é um exemplo de Infraestrutura como Serviço (IaaS). IaaS é um modelo de nuvem em que os recursos de infraestrutura, como servidores, armazenamento e rede, são fornecidos como um serviço sob demanda. Ao executar o PostgreSQL em uma VM do Azure, você está usando a infraestrutura do Azure para hospedar seu banco de dados.

upvoted 1 times

🗨️ 👤 **Kanatius** 1 year, 4 months ago

Sim, IaaS e PaaS. Li a doc agora há pouco sobre SQL server nas VMs dizendo que era IaaS, o postgres não seria diferente

upvoted 1 times

🗨️ 👤 **[Removed]** 1 year, 4 months ago

IaaS and PaaS are the correct answer.

The solution is wrong so bear that in mind.

upvoted 2 times

Which Azure storage solution provides native support for POSIX-compliant access control lists (ACLs)?

- A. Azure Table storage
- B. Azure Data Lake Storage
- C. Azure Queue storage
- D. Azure Files

**Suggested Answer:** B

Reference:

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

Community vote distribution

B (100%)

- 🗨️ **yusufuthman57** Highly Voted 2 years, 10 months ago  
Azure Data Lake Storage Gen2 is the correct Answer.  
upvoted 36 times
- 🗨️ **RichardWilde** Highly Voted 2 years, 10 months ago  
Azure Data Lake Storage  
upvoted 13 times
- 🗨️ **TKC\_2023** Most Recent 1 month, 1 week ago  
Selected Answer: B  
Azure Data Lake Storage Gen2  
upvoted 1 times
- 🗨️ **AGTraining** 6 months, 2 weeks ago  
Azure Data Lake Storage  
upvoted 1 times
- 🗨️ **erwin\_t82** 10 months, 3 weeks ago  
B answer is correct - Azure Data Lake Storage.  
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).  
<https://learn.microsoft.com/en-us/azure/storage/blobs/data-lake-storage-access-control>  
upvoted 3 times
- 🗨️ **TKC\_2023** 11 months, 2 weeks ago  
Selected Answer: B  
POSIX-ADLS  
upvoted 1 times
- 🗨️ **DaliborUhlir** 1 year, 1 month ago  
to memorize this ADLS ~ ACLs  
upvoted 8 times
- 🗨️ **akhil5432** 1 year, 1 month ago  
AZURE DATA LAKE STORAGE  
upvoted 1 times
- 🗨️ **Sender** 1 year ago  
ALRIGHT THANKS BRO  
upvoted 3 times
- 🗨️ **chiiiweiii** 1 year, 6 months ago  
Selected Answer: B  
POSIX-ADLS  
upvoted 4 times

🗨️ **ravick4u** 1 year, 7 months ago

**Selected Answer: B**

Azure Data Lake Storage

upvoted 2 times

🗨️ **roop777** 1 year, 8 months ago

Azure Data Lake Storage

upvoted 2 times

🗨️ **ria97** 2 years, 1 month ago

Was in exam 11/08. Correct answer

upvoted 5 times

🗨️ **Eltooth** 2 years, 2 months ago

**Selected Answer: B**

B is correct answer.

upvoted 5 times

🗨️ **RichMelbourne** 2 years, 5 months ago

Shown on test 15/04/2022

upvoted 6 times

🗨️ **Baha2021** 2 years, 5 months ago

**Selected Answer: B**

Answer is B

upvoted 3 times

🗨️ **SiDoCiOuS** 2 years, 5 months ago

**Selected Answer: B**

Appeared on the exam on 03/22/2022

Answer is correct!

upvoted 4 times

🗨️ **Koph** 2 years, 9 months ago

Azure data lake storage is correct

upvoted 3 times

HOTSPOT -

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

Hot Area:

### 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.

### Answer Area

Relational data is stored in

Suggested Answer:

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

 **ahee** Highly Voted 2 years, 10 months ago

Aswer is C.

A relational database stores data in tables. Tables are organized into columns and row.

upvoted 52 times

 **mamahani** Highly Voted 1 year, 5 months ago

I took the exam 12.04.2023 and got 1000, 100% score; this question appeared, correct answer:

a tabular form of rows and columns

upvoted 7 times

 **heobanten04** Most Recent 5 months, 2 weeks ago

took the exam on 30/03/2024 and this question appeared. Pick C - a tabular form of rows and columns

upvoted 2 times

 **AGTraining** 6 months, 2 weeks ago

Answer is C.

upvoted 1 times

 **Edgy\_San** 7 months, 1 week ago

@DSilisbury In the context of relational databases, "tabular" refers to the structure of the data representation. Data is organized in tables, which are two-dimensional structures consisting of rows and columns. Each row in the table represents a record or a specific instance of data, and each column represents a different attribute or field of the data.

Tabular means related to or arranged in a table. Spreadsheets do have tabs. But the questions is inquiring about tabular and refers to a table. (singular)

The term "tabular" is used to describe the tabular form of organization, emphasizing the arrangement of data into a grid-like structure with rows and columns, similar to a table. This structure is a fundamental characteristic of relational databases, providing a systematic and organized way to store, retrieve, and manage data.

upvoted 2 times

 **BXtheCoder** 7 months, 3 weeks ago

Relational data is stored in columns and rows. The answer is correct.

upvoted 1 times

- 🗨️ 👤 **fe6d44f** 8 months, 1 week ago  
Passed exam on 9/1/2024, got 984/1000.  
My answer was a tabular form of rows and columns  
upvoted 1 times
- 🗨️ 👤 **MannyJZ** 1 year, 1 month ago  
On exam today 08/08/23  
upvoted 2 times
- 🗨️ 👤 **VNR25** 1 year, 2 months ago  
This question appeared in the 07/07/23 exam and the answer is correct.  
upvoted 1 times
- 🗨️ 👤 **Dudures** 1 year, 5 months ago  
Exam 04/15/2023  
upvoted 3 times
- 🗨️ 👤 **ahmedkmj** 1 year, 5 months ago  
in the exam 26/03  
upvoted 1 times
- 🗨️ 👤 **PhyoPluto** 1 year, 6 months ago  
the exam on 07/03/2023  
upvoted 1 times
- 🗨️ 👤 **chiiiweiii** 1 year, 6 months ago  
Correct Answer  
upvoted 2 times
- 🗨️ 👤 **Nadjib** 1 year, 8 months ago  
Aswer is C. A relational database stores data in tables.  
upvoted 1 times
- 🗨️ 👤 **edurakhan** 1 year, 10 months ago  
In exam today 11/15/22  
upvoted 1 times
- 🗨️ 👤 **Albertopology** 1 year, 11 months ago  
on exam 10 Oct 2022  
upvoted 3 times
- 🗨️ 👤 **blackskies** 1 year, 11 months ago  
was on exam today 09/30/22  
upvoted 3 times

HOTSPOT -

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

Hot Area:

### Answer Area

A block of code that runs in a database is called

<input type="text"/>
a stored procedure.
a table.
a view.
an index.

### Answer Area

A block of code that runs in a database is called

Suggested Answer:

<input type="text"/>
a stored procedure.
a table.
a view.
an index.

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/stored-procedures/stored-procedures-database-engine?view=sql-server-ver15>

 **Marksmen** Highly Voted 2 years, 10 months ago

Stored Procedure is correct  
upvoted 51 times

 **yusufuthman57** Highly Voted 2 years, 10 months ago

Stored Procedure!  
upvoted 14 times

 **AGTraining** Most Recent 6 months, 2 weeks ago

Stored Procedure  
upvoted 1 times

 **BXtheCoder** 7 months, 3 weeks ago

Yep. It's a Stored Procedure. A stored procedure stores a command or even a line of code that can be reused in the future.  
upvoted 2 times

 **trashbox** 10 months, 2 weeks ago

A.

Appeared on 2023-11-05.

upvoted 1 times

 **VNR25** 1 year, 2 months ago

This question appeared in the 07/07/23 exam and the answer is correct.  
upvoted 1 times

 **anilstartbit** 1 year, 3 months ago

SP is correct  
upvoted 1 times

 **PhyoPluto** 1 year, 6 months ago

the exam on 07/03/2023  
upvoted 4 times

🗨️ 👤 **Kirron** 2 years, 1 month ago

Stored procedure is correct  
upvoted 5 times

🗨️ 👤 **Eltooth** 2 years, 2 months ago

Stored procedure is correct answer.  
upvoted 3 times

🗨️ 👤 **Deliboi** 2 years, 2 months ago

The rest dont make sense so it has to be Stored Procedure  
upvoted 5 times

🗨️ 👤 **meldr** 2 years, 5 months ago

Yes, stored procedure  
upvoted 2 times

🗨️ 👤 **MCOMP** 2 years, 6 months ago

In terms of view as well, the code can be in a block i.e. transaction  
upvoted 4 times

🗨️ 👤 **Grsha** 2 years, 7 months ago

Stored Procedure  
upvoted 3 times

🗨️ 👤 **Haidardba** 2 years, 9 months ago

Stored Procedure is Correct Answer.  
upvoted 5 times

DRAG DROP -

Match the types of analytics that can be used to answer the business questions.

To answer, drag the appropriate analytics type from the column on the left to its question on the right. Each analytics type may be used once, more than once, or not at all.

NOTE: Each correct match is worth one point.

Select and Place:

### Analytics Types

### Answer Area

Cognitive	Why did sales increase last month?	<input type="text"/>
Diagnostic	Provide specific guidance on how to allocate current budget to buy different inventory items.	<input type="text"/>
Descriptive	Detect images of the company's logo included in PDF documents	<input type="text"/>
Predictive		
Prescriptive		

**Suggested Answer:** *Incorrect Answer:*

Box 1: Diagnostic -

Diagnostic Analytics: At this stage you can begin to answer some of those why questions. Historical data can begin to be measured against other data to answer the question of why something happened in the past. This is the process of gathering and interpreting different data sets to identify anomalies, detect patterns, and determine relationships.

Box 2: Prescriptive -

Prescriptive analytics is a combination of data, mathematical models, and various business rules to infer actions to influence future desired outcomes.

Predictive analytics, broadly speaking, is a category of business intelligence that uses descriptive and predictive variables from the past to analyze and identify the likelihood of an unknown future outcome

Box 3: Cognitive -

Reference:

<https://demand-planning.com/2020/01/20/the-differences-between-descriptive-diagnostic-predictive-cognitive-analytics/>

<https://azure.microsoft.com/en-us/blog/answering-whats-happening-whys-happening-and-what-will-happen-with-iot-analytics/>

 **Afetik** Highly Voted 2 years, 4 months ago

The answers for this question are correct.  
upvoted 49 times

 **Eltooth** Highly Voted 2 years, 2 months ago

Diagnostic

Predictive

Cognitive

upvoted 21 times

 **Eltooth** 2 years, 2 months ago

Typo!

Diag

Prescriptive

Cognitive

upvoted 37 times

 **profesorklaus** Most Recent 1 week, 2 days ago

Correct

upvoted 1 times

🗨️ 👤 **\_Vijay\_** 1 month, 3 weeks ago

The answers are correct.

upvoted 1 times

🗨️ 👤 **AGTraining** 6 months, 2 weeks ago

Diagnostic

Prescriptive

Cognitive

upvoted 2 times

🗨️ 👤 **chiiiweiii** 1 year, 6 months ago

Correct Answer

upvoted 4 times

🗨️ 👤 **necole27** 1 year, 9 months ago

correct

upvoted 3 times

🗨️ 👤 **Hadley205** 2 years, 2 months ago

I think this question is the same as Q6? So the last statement should be Cognitive, not Descriptive as in Q6

upvoted 5 times

🗨️ 👤 **szl0144** 2 years, 2 months ago

The keyword is detects, cognitive is for detections. In #6, the question asks "which thing". It is an objective instead of a detection action.

upvoted 7 times

🗨️ 👤 **Deliboi** 2 years, 2 months ago

They are correct

upvoted 4 times

🗨️ 👤 **Fatima\_Dashdamirova** 2 years, 3 months ago

correct aswer

upvoted 6 times

🗨️ 👤 **Duncan567710** 2 years, 4 months ago

Correct

upvoted 5 times

🗨️ 👤 **bfan95** 2 years, 4 months ago

Correct

upvoted 4 times

🗨️ 👤 **trdr** 2 years, 4 months ago

Correct Answer!

upvoted 3 times

Which type of database is Azure Database for PostgreSQL?

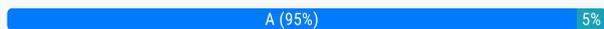
- A. Platform as a service (PaaS)
- B. Infrastructure as a service (IaaS)
- C. Microsoft SQL Server
- D. on-premises

**Suggested Answer: A**

Reference:

<https://docs.microsoft.com/en-us/azure/postgresql/overview-postgres-choose-server-options>

Community vote distribution



- 🗨️ **bfan95** Highly Voted 2 years, 4 months ago  
Correct answer: A, PAAS  
upvoted 32 times
- 🗨️ **Afetik** Highly Voted 2 years, 4 months ago  
The answer is correct - A  
upvoted 14 times
- 🗨️ **TKC\_2023** Most Recent 1 month, 1 week ago  
Selected Answer: A  
PaaS is the answer  
upvoted 1 times
- 🗨️ **\_Vijay\_** 1 month, 3 weeks ago  
Selected Answer: A  
Yep! It is correct.  
upvoted 1 times
- 🗨️ **okayhey** 6 months, 1 week ago  
All database services are PaaS  
upvoted 1 times
- 🗨️ **AGTraining** 6 months, 2 weeks ago  
Option A  
upvoted 1 times
- 🗨️ **BXtheCoder** 7 months, 3 weeks ago  
The answer is A, PaaS. Azure Databases uses DBaaS, which falls under PaaS.  
upvoted 1 times
- 🗨️ **fe6d44f** 8 months, 1 week ago  
Passed exam on 9/1/2024, got 984/1000.  
My answer was PAAS  
upvoted 4 times
- 🗨️ **ravinperera** 9 months, 3 weeks ago  
Azure Database for PostgreSQL is falls under the category of Database as a Service (DBaaS) or Platform as a Service (PaaS).  
upvoted 2 times
- 🗨️ **AZFabio** 1 year ago  
Selected Answer: A  
All DB offerings in Azure are PaaS, except when installed in a VM (then it is IaaS).  
upvoted 9 times
- 🗨️ **MannyJZ** 1 year, 1 month ago  
On exam today 08/08/23

upvoted 2 times

🗨️ 👤 **rohirag** 1 year, 2 months ago

Selected Answer: A

PaaS.

<https://learn.microsoft.com/en-us/training/modules/explore-provision-deploy-relational-database-offerings-azure/3-azure-database-open-source>

upvoted 1 times

🗨️ 👤 **chematab** 1 year, 4 months ago

Selected Answer: A

Correct Answer

upvoted 1 times

🗨️ 👤 **chiiiweiii** 1 year, 6 months ago

Selected Answer: A

Correct Answer

upvoted 2 times

🗨️ 👤 **tangerine141** 1 year, 6 months ago

Selected Answer: A

Correct answer: A, PAAS

upvoted 1 times

🗨️ 👤 **Joanna0** 1 year, 7 months ago

Selected Answer: A

PaaS is the correct answer.

upvoted 1 times

🗨️ 👤 **ravick4u** 1 year, 7 months ago

Selected Answer: A

Platform as a service

upvoted 1 times

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

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

**Suggested Answer: A**

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).

Reference:

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

Community vote distribution

A (100%)

 **svetlanam** Highly Voted 2 years, 4 months ago

Correct Answer is A. Azure Data Lake Storage  
upvoted 25 times

 **TKC\_2023** Most Recent 1 month, 1 week ago

Selected Answer: A

Correct Answer is A. Azure Data Lake Storage  
upvoted 1 times

 **AGTraining** 6 months, 2 weeks ago

Option A  
upvoted 1 times

 **AZFabio** 1 year ago

Selected Answer: A

correct  
upvoted 2 times

 **srgopalkr** 1 year, 1 month ago

This is a trick question as both Azure Blob Storage & Azure Data Lake support ACL's at the file and directory level leveraging Security principal (user, group, service principal, or managed identity). Additionally, Azure Data Lake storage system is based on Azure Blob Storage system. However Azure Data Lake Gen2, implements ACL differently in comparison to Azure Blob Storage, with the former supporting Azure RBAC & POSIX ACL's. Azure Blob Storage implements access control using SAS and AAD authentication. So in my opinion, both A & C are correct.  
upvoted 3 times

 **HRKK** 1 year, 2 months ago

CORRECT ANSWER A  
upvoted 1 times

 **ahmedkmj** 1 year, 5 months ago

in the exam 26/03  
upvoted 3 times

 **chiiweiii** 1 year, 6 months ago

Correct Answer  
upvoted 2 times

 **Joanna0** 1 year, 7 months ago

Selected Answer: A

<https://learn.microsoft.com/en-us/azure/storage/blobs/data-lake-storage-access-control#about-acls>  
upvoted 1 times

🗨️ **ravick4u** 1 year, 7 months ago

**Selected Answer: A**

A. Azure Data Lake Storage

upvoted 1 times

🗨️ **nilsoncarlosf** 1 year, 9 months ago

Correct is A

upvoted 1 times

🗨️ **JA2018** 1 year, 10 months ago

**Selected Answer: A**

Agreed, A is the correct option. duplicated question

upvoted 2 times

🗨️ **Albertopology** 1 year, 11 months ago

on exam 10 Oct 2022

upvoted 3 times

🗨️ **ria97** 2 years, 1 month ago

Was in exam 11/08. Correct answer

upvoted 3 times

🗨️ **Breeze21** 2 years, 2 months ago

Exam: 19.07.22

upvoted 3 times

🗨️ **Eltooth** 2 years, 2 months ago

**Selected Answer: A**

A is correct answer. ADL

upvoted 2 times

🗨️ **simplementeluca** 2 years, 2 months ago

**Selected Answer: A**

A the correct

upvoted 1 times

HOTSPOT -

Select the answer that correctly completes the sentence.

Hot Area:

### Answer Area

A  is a chart of colored rectangles. The size of the rectangles represent

line chart
matrix
scatter
treemap

the relative value of each item. The chart can be hierarchical, displaying data as a set of rectangles nested within the main rectangle.

Suggested Answer:

### Answer Area

A  is a chart of colored rectangles. The size of the rectangles represent

line chart
matrix
scatter
treemap

the relative value of each item. The chart can be hierarchical, displaying data as a set of rectangles nested within the main rectangle.

A treemap chart divides the chart area into rectangles that represent the different levels and relative sizes of the data hierarchy.

Reference:

<https://docs.microsoft.com/en-us/sql/reporting-services/report-design/tree-map-and-sunburst-charts-in-reporting-services?view=sql-server-ver15>

 **bfan95** Highly Voted 2 years, 4 months ago

treemap

upvoted 30 times

 **mamahani** Highly Voted 1 year, 5 months ago

I took the exam 12.04.2023 and got 1000, 100% score; this question appeared, correct answer:

treemap

upvoted 6 times

 **heobanten04** Most Recent 5 months, 2 weeks ago

I took the exam 30.03.2024 and got 936; this question appeared, correct answer:

treemap

upvoted 2 times

 **Tosceen** 1 month, 2 weeks ago

I hope the dumps here is enough to pass the exam

upvoted 2 times

 **AGTraining** 6 months, 2 weeks ago

treemap

upvoted 1 times

 **BXtheCoder** 7 months, 3 weeks ago

Yeah. The answer is Treemap.

upvoted 1 times

- 🗨️ 👤 **james2033** 10 months, 3 weeks ago  
treemap - colored rectangles, can be hierarchical.  
upvoted 1 times
- 🗨️ 👤 **HRKK** 1 year, 2 months ago  
CORRECT  
upvoted 1 times
- 🗨️ 👤 **chiiiweiii** 1 year, 6 months ago  
Correct Answer  
upvoted 2 times
- 🗨️ 👤 **jhamakalan** 1 year, 7 months ago  
answer should be C option  
upvoted 1 times
- 🗨️ 👤 **Alex4190** 1 year, 6 months ago  
Wrong, it is treemap  
upvoted 1 times
- 🗨️ 👤 **BXtheCoder** 7 months, 3 weeks ago  
No. C is a Scatter Plot which is used to analyze what happens to a variable when another variable changes, consist of many dots on the chart.  
Treemap consists of coloured rectangles.  
upvoted 1 times
- 🗨️ 👤 **blackskies** 1 year, 11 months ago  
was on exam today 09/30/22  
upvoted 4 times
- 🗨️ 👤 **gabrisiq** 2 years, 1 month ago  
Colored Rectangles = Treemap  
upvoted 4 times
- 🗨️ 👤 **Eltooth** 2 years, 2 months ago  
Tree map is correct answer.  
upvoted 2 times
- 🗨️ 👤 **WWDITDA** 2 years, 2 months ago  
Treemap - Correct.  
upvoted 1 times
- 🗨️ 👤 **Manoj6** 2 years, 3 months ago  
Treemap is the correct answer  
upvoted 2 times
- 🗨️ 👤 **dk87** 2 years, 4 months ago  
Exam: 05/18/2022  
upvoted 3 times

What is a characteristic of batch processing?

- A. The data ingested during batch processing must be processed as soon as the data is received.
- B. Large datasets must be split into batches of less than 1 GB before the data can be processed.
- C. There is a noticeable time delay between ingesting data and obtaining the data processing results.
- D. Batch processing can only process data that is structured.

**Suggested Answer:** D

Community vote distribution

C (99%)

🗳️ **sallymaher** Highly Voted 2 years, 5 months ago

Answer should be C  
upvoted 141 times

🗳️ **ha1p** 1 year, 5 months ago

agree w/ you  
upvoted 1 times

🗳️ **JoJoY** Highly Voted 2 years, 4 months ago

Selected Answer: C

The answer is C  
upvoted 44 times

🗳️ **skaha** 2 years, 4 months ago

which is not characteristic of Batch processing>>>Then the answer is D  
upvoted 4 times

🗳️ **profesorklaus** Most Recent 1 week, 2 days ago

Selected Answer: C

C is correct answer  
upvoted 1 times

🗳️ **\_Vijay\_** 1 month, 3 weeks ago

Selected Answer: C

The "D" answer is a BIG time wrong answer.  
upvoted 1 times

🗳️ **Jikkuz** 3 months, 4 weeks ago

Selected Answer: C

C is the answer  
upvoted 1 times

🗳️ **heobanten04** 5 months, 2 weeks ago

I took the exam 30.03.2024 and got 936; this question appeared, correct answer: C  
upvoted 5 times

🗳️ **IK13** 5 months, 3 weeks ago

C

<https://learn.microsoft.com/en-us/training/modules/explore-fundamentals-stream-processing/2-batch-stream>  
upvoted 2 times

🗳️ **AGTraining** 6 months, 2 weeks ago

option C  
upvoted 1 times

🗳️ **Edgy\_San** 7 months, 1 week ago

The way I understand batch processing it has to either way for an event to occur or a condition to be met to be separated the batches of data.  
Example: End of the month credit card bill goes out.

upvoted 1 times

🗨️ 👤 **Edgy\_San** 7 months, 1 week ago

\*wait (Typo up there)

upvoted 1 times

🗨️ 👤 **masadamsahid** 10 months, 3 weeks ago

**Selected Answer: C**

Choice D should not be correct because batch processing can process semi-structured (NoSQL) and even unstructured data (with complex and more effort) too.

The most logical answer is C, it's expected there will be a delay.

upvoted 3 times

🗨️ 👤 **james2033** 10 months, 3 weeks ago

**Selected Answer: C**

keyword 'delay' for a batch.

upvoted 3 times

🗨️ 👤 **ig\_85** 11 months, 1 week ago

**Selected Answer: C**

The fundamental requirement of such batch processing engines is to scale out computations to handle a large volume of data. Unlike real-time processing, batch processing is expected to have latencies

upvoted 1 times

🗨️ 👤 **t201** 11 months, 2 weeks ago

Answer is C

upvoted 1 times

🗨️ 👤 **Yuezhen** 11 months, 4 weeks ago

**Selected Answer: C**

C should be the correct answer.

upvoted 2 times

🗨️ 👤 **Ratz** 1 year ago

Option C :

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

Read this para:

The fundamental requirement of such batch processing engines is to scale out computations to handle a large volume of data. Unlike real-time processing, batch processing is expected to have latencies (the time between data ingestion and computing a result) that measure in minutes to hours.

upvoted 2 times

🗨️ 👤 **Ifparauj0** 1 year ago

**Selected Answer: C**

Yes, it is possible to use batch processing with unstructured data. Batch processing is a data processing technique that involves collecting and processing data in predefined batches or groups. It can be used with both structured and unstructured data, although the approach and tools used may differ.

upvoted 1 times

🗨️ 👤 **AZFabio** 1 year ago

**Selected Answer: C**

C is correct

upvoted 1 times

HOTSPOT -

Select the answer that correctly completes the sentence.

Hot Area:

### Answer Area

In a SQL statement,  are examples of aggregate functions.

- FROM and WHERE
- GROUP BY and ORDER BY
- JOIN and MERGE
- MAX and MIN

Suggested Answer:

### Answer Area

In a SQL statement,  are examples of aggregate functions.

- FROM and WHERE
- GROUP BY and ORDER BY
- JOIN and MERGE
- MAX and MIN

An aggregate function performs a calculation on a set of values, and returns a single value. The following are aggregate functions:

- ⇒ APPROX\_COUNT\_DISTINCT
- ⇒ AVG
- ⇒ CHECKSUM\_AGG
- ⇒ COUNT
- ⇒ COUNT\_BIG

GROUPING -

- 
- ⇒ GROUPING\_ID
- ⇒ MAX
- ⇒ MIN
- ⇒ STDEV
- ⇒ STDEVP
- ⇒ STRING\_AGG
- ⇒ SUM
- ⇒ VAR
- ⇒ VARP

Reference:

<https://docs.microsoft.com/en-us/sql/t-sql/functions/aggregate-functions-transact-sql?view=sql-server-ver15>

 **Eltooth** Highly Voted 2 years, 2 months ago

Max and Min

upvoted 18 times

 **mamahani** Highly Voted 1 year, 5 months ago

I took the exam 12.04.2023 and got 1000, 100% score; this question appeared, correct answer:

MAX and MIN

upvoted 14 times

 **heobanten04** Most Recent 5 months, 2 weeks ago

I took the exam on 30.03.2024 and got 936; this question appeared, correct answer: MAX and MIN

upvoted 2 times

🗨️ **AGTraining** 6 months, 2 weeks ago

max and min

upvoted 1 times

🗨️ **james2033** 10 months, 3 weeks ago

MAX and MIN are aggregate functions.

upvoted 1 times

🗨️ **chiiiweiii** 1 year, 6 months ago

Correct Answer

upvoted 3 times

🗨️ **YOGESHGP19111982** 1 year, 7 months ago

MIN MAZ

upvoted 1 times

🗨️ **Abains** 1 year, 10 months ago

Correct

upvoted 1 times

🗨️ **dk87** 2 years, 4 months ago

Exam: 05/18/2022

upvoted 5 times

🗨️ **Ashishsingh07** 2 years, 4 months ago

MIN MAX

upvoted 3 times

🗨️ **bfan95** 2 years, 4 months ago

correct

upvoted 4 times

🗨️ **arif17** 2 years, 4 months ago

answer is correct.

upvoted 3 times

🗨️ **pads98** 2 years, 4 months ago

answer is correct.

upvoted 4 times

HOTSPOT -

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE: Each correct selection is worth one point.

Hot Area:

### Answer Area

Statements	Yes	No
A job that processes sales data once daily is an example of a batch workload.	<input type="radio"/>	<input type="radio"/>
A job that calculates a rolling average temperature reading is an example of a streaming workload.	<input type="radio"/>	<input type="radio"/>
A job that calculates average revenue per product for the last month is an example of a streaming workload.	<input type="radio"/>	<input type="radio"/>

Suggested Answer:

### Answer Area

Statements	Yes	No
A job that processes sales data once daily is an example of a batch workload.	<input checked="" type="radio"/>	<input type="radio"/>
A job that calculates a rolling average temperature reading is an example of a streaming workload.	<input checked="" type="radio"/>	<input type="radio"/>
A job that calculates average revenue per product for the last month is an example of a streaming workload.	<input type="radio"/>	<input checked="" type="radio"/>

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 -

Reference:

<https://www.geeksforgeeks.org/difference-between-batch-processing-and-stream-processing/>

 **Edriv** Highly Voted 1 year, 9 months ago

Yes, Yes, No

upvoted 39 times

 **heenak55** Highly Voted 2 years, 4 months ago

In exam 21 April 2022

upvoted 16 times

 **b9f0525** Most Recent 3 months, 2 weeks ago

Yes, Yes, No

upvoted 1 times

 **AGTraining** 6 months, 2 weeks ago

Yes, Yes, No

upvoted 1 times

 **Edgy\_San** 7 months, 1 week ago

I thought it was yes on the second statement and according to chatGPT here is the scoop for your consideration: "Yes, a job that calculates a rolling average of temperature readings can be considered an example of a streamed workload. In this context, a streaming workload typically involves processing and analyzing data in real-time or near-real-time as it is generated."

When dealing with temperature readings, a streaming workload might involve continuously receiving and processing data from sensors or sources providing temperature information. Calculating a rolling average in real-time allows for dynamic and up-to-date insights into temperature trends without the need to wait for batch processing.

So, in summary, calculating a rolling average for temperature readings aligns with the characteristics of a streaming workload, especially if the analysis occurs as new data is continuously generated.

upvoted 1 times

🗨️ 👤 **victorhugo13** 10 months, 2 weeks ago

yes yes, no

upvoted 1 times

🗨️ 👤 **akhil5432** 1 year, 1 month ago

yes yes no

upvoted 1 times

🗨️ 👤 **DataEngDP** 1 year, 4 months ago

The second one is indeed streaming, upon searching on it Microsoft has a good resource that proves that an average calculation that is done over a period of time uses Streaming analytics, for example a taxi company that wants to know average tip on a 5 minutes window. In this case the rolling avg temperature it is processed as it arrives which is why is streaming. There is NO waiting time in this case.

<https://learn.microsoft.com/en-us/azure/architecture/reference-architectures/data/stream-processing-stream-analytics>

upvoted 2 times

🗨️ 👤 **DataEngDP** 1 year, 4 months ago

Rolling means in this case that new info about the temperatures is updated as it arrives.

upvoted 1 times

🗨️ 👤 **ahmedkmj** 1 year, 5 months ago

in the exam 26/03

upvoted 4 times

🗨️ 👤 **chiiiweiii** 1 year, 6 months ago

YYN are Correct Answer

upvoted 4 times

🗨️ 👤 **A\_Fardeen** 1 year, 6 months ago

it's Y,Y,N

upvoted 3 times

🗨️ 👤 **noussa** 2 years, 1 month ago

Yes, Yes( rolling), No

upvoted 8 times

🗨️ 👤 **gabrisiq** 2 years, 1 month ago

Batch, Streaming, Batch

upvoted 5 times

🗨️ 👤 **kuiberbelt** 2 years, 1 month ago

A rolling or moving average is a statistic that captures the average change in a data series over time.

upvoted 2 times

🗨️ 👤 **xRiot007** 2 years, 2 months ago

Yes - periodical jobs use batching

Yes - a rolling average will use a sized set that is updated with newest data.

No - Last month sales are past data only, so this is a batch job.

upvoted 7 times

🗨️ 👤 **Eltooth** 2 years, 2 months ago

Yes

No - average is taken then processed over a period of time -> batch

No

upvoted 7 times

  **schurick13** 2 years, 2 months ago

The key word is "ROLLING average". Hence stream processing

upvoted 8 times

  **dk87** 2 years, 4 months ago

Exam: 05/18/2022

upvoted 4 times

HOTSPOT -

Select the answer that correctly completes the sentence.

Hot Area:

### Answer Area

In a relational database, each row in a table has

	▼
a different set of columns	
a key-value pair	
the same set of columns	
unstructured data	

Suggested Answer:

### Answer Area

In a relational database, each row in a table has

	▼
a different set of columns	
a key-value pair	
the same set of columns	
unstructured data	

The same set of columns.

In relational databases, a row is a data record within a table. Each row, which represents a complete record of specific item data, holds different data within the same structure. A row is occasionally referred to as a tuple.

Incorrect:

Not: a key value pair.

Unlike relational databases, key-value databases do not have a specified structure. Relational databases store data in tables where each column has an assigned data type. Key-value databases are a collection of key-value pairs that are stored as individual records and do not have a predefined data structure.

Reference:

<https://www.techopedia.com/definition/4425/database-row>

<https://www.techtarget.com/searchdatamanagement/tip/NoSQL-database-types-explained-Key-value-store>

DBA\_newbie Highly Voted 2 years ago

correct

upvoted 28 times

Albertology Highly Voted 1 year, 11 months ago

on exam 10 Oct 2022

upvoted 7 times

AGTraining Most Recent 6 months, 2 weeks ago

the same set of columns

upvoted 1 times

fe6d44f 8 months, 1 week ago

Passed exam on 9/1/2024, got 984/1000.

My answer was the same set of columns

upvoted 1 times

james2033 10 months, 3 weeks ago

In a relational database, each row in a table has THE SAME SET OF COLUMNS.

upvoted 1 times

🗨️ 👤 **globy118** 1 year, 2 months ago

Was on exam 24/6/2023

upvoted 4 times

🗨️ 👤 **ahmedkmj** 1 year, 5 months ago

in the exam 26/03

upvoted 4 times

🗨️ 👤 **Kseetha** 1 year, 6 months ago

Ans: Same set of columns

upvoted 3 times

🗨️ 👤 **chiiweiii** 1 year, 6 months ago

Correct Answer

upvoted 2 times

🗨️ 👤 **PatYeo** 1 year, 11 months ago

correct

upvoted 6 times

HOTSPOT -

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE: Each correct selection is worth one point.

Hot Area:

### Answer Area

Statements	Yes	No
Normalizing a database increases the throughput of writing transactions.	<input type="radio"/>	<input type="radio"/>
Analytics systems are more normalized than transactional systems.	<input type="radio"/>	<input type="radio"/>
Normalizing a database results in queries that require more joins.	<input type="radio"/>	<input type="radio"/>

### Suggested Answer:

#### Answer Area

Statements	Yes	No
Normalizing a database increases the throughput of writing transactions.	<input type="radio"/>	<input checked="" type="radio"/>
Analytics systems are more normalized than transactional systems.	<input type="radio"/>	<input checked="" type="radio"/>
Normalizing a database results in queries that require more joins.	<input checked="" type="radio"/>	<input type="radio"/>

Box 1: No -

Database normalization is the process of restructuring a relational database in accordance with a series of so-called normal forms in order to reduce data redundancy and improve data integrity.

Full normalisation will generally not improve performance, in fact it can often make it worse but it will keep your data duplicate free.

Box 2: No -

Analytics systems are denormalized to increase performance.

Transactional database systems are normalized to increase data consistency.

Box 3: Yes -

Transactional database systems are more normalized and requires more joins.

Reference:

<https://www.sqlshack.com/what-is-database-normalization-in-sql-server>

 **Kris\_ETL** Highly Voted 2 years ago

YES, NO, YES

Normalizing decreases throughput for reading (need to make joins), but increases for writing (Less to write, since there are no duplicates)

upvoted 108 times

 **\_Vijay\_** 1 month, 2 weeks ago

This is not correct. Answer is NO, NO, YES. Normalizing will involve foreign keys, thus more checks when inserting records, hence it will be slow down write operation.

upvoted 1 times

  **nvravan** 1 year, 8 months ago

But, any foreign key constraint may slow down the throughput of the write operation.

upvoted 5 times

  **okayhey** 6 months, 1 week ago

A foreign key is a direct reference, so no. Having to go through all data takes more time. I agree with Kris\_ETL

upvoted 2 times

  **NareshKothari** 1 year, 11 months ago

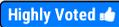
Agree. Less Write means better performance. Yes, No, Yes is the right answer.

upvoted 10 times

  **Shahbaz21** 1 year, 5 months ago

Agreed

upvoted 2 times

  **Jiviify**  2 years ago

NO

NO

YES

Answers are correct

Normalizing a Database increases the throughput for reading but reduces for writing.

A normalized database is advantageous when operations are write-intensive or ACID compliance is required. Some advantages include:

Updates run quickly due to no data being duplicated in multiple locations.

Inserts run quickly since there is only a single insertion point for a piece of data, and no duplication is required.

Tables are typically smaller than the tables found in non-normalized databases. This usually allows the tables to fit into the buffer, thus offering faster performance.

Data integrity and consistency are an absolute must if the database must be ACID compliant. A normalized database helps immensely with such an undertaking.

upvoted 28 times

  **raydel92** 1 year, 11 months ago

The first one is a YES. Your answer contradicts what you develop later. Source: <https://dzone.com/articles/pros-and-cons-of-database-normalization>

upvoted 13 times

  **profesorklaus**  1 week, 2 days ago

The first question is silly and non-obvious. Database normalization can improve write throughput, but not always. It depends on the complexity of the database it is applied to. If there are many foreign keys and the system is updating many tables at once, then no. If there are mostly write operations on single tables, then yes. In my experience as a developer, we normalize to improve write and read throughput, but in some rare circumstances we do not. I would select YES for the first checkbox.

upvoted 1 times

  **TKC\_2023** 1 month, 1 week ago

Normalizing a database typically involves organizing the data to reduce redundancy and improve data integrity. While normalization can make the database more efficient for read operations and reduce data anomalies, it often leads to more complex data structures with multiple related tables. As a result, writing transactions can become more complex and slower because they may require more operations, such as inserting data into multiple tables and maintaining foreign key relationships.

Thus, normalization generally decreases the throughput of writing transactions rather than increasing it.

upvoted 1 times

  **Aqshata** 2 months, 2 weeks ago

Yes, No, Yes

upvoted 1 times

  **Sekharad** 3 months, 3 weeks ago

1 > False.

Normalizing a database generally involves organizing the data to reduce redundancy and improve data integrity. While this can make read operations more efficient by minimizing data duplication and ensuring consistency, it often has the opposite effect on write operations.

upvoted 1 times

🗨️ **AGTraining** 6 months, 2 weeks ago

Si, NO, Si

upvoted 1 times

🗨️ **LM12** 6 months, 4 weeks ago

According to Chatgpt: Normalized databases are typically better optimized for write operations (INSERT, UPDATE, DELETE) rather than read operations. The normalization process involves organizing data to minimize redundancy and dependency, which can improve data integrity but may lead to more complex query structures for certain read operations.

upvoted 1 times

🗨️ **Yomzie** 7 months ago

Take it from a DBA with over 20+ years of work experience:

Yes: Normalized systems are optimized for WRITE ops; much less so for READ ops (due to join operation btw different tables).

No: Analytic systems are less normalized; they have more AGGREGATE data.

Yes: Queries on normalized tables require more TABLE JOINS (full, inner, left, right... joins)

upvoted 5 times

🗨️ **james2033** 10 months, 3 weeks ago

Normalizing a database increases the throughput of writing transactions. --> No.

Analytics systems are more normalized than transactional systems. --> No.

Normalizing a database results in queries that require more joins. --> Yes.

upvoted 2 times

🗨️ **hcq31818** 11 months, 4 weeks ago

I think the first should be No as

Normalizing a database typically improves data integrity and reduces data redundancy but can sometimes lead to a slight decrease in write transaction throughput. This is because, in a normalized database, data is distributed across multiple tables, and writing transactions may require inserting or updating records in multiple tables, which can add some overhead compared to denormalized databases. However, the trade-off is improved data consistency and easier maintenance.

Second - No

Third - Yes

upvoted 2 times

🗨️ **HRKK** 1 year, 2 months ago

NO

NO

YES

CORRECT ONE

upvoted 3 times

🗨️ **Vanessa23** 1 year, 5 months ago

Yes-No-Yes. Writing is optimized for normalized databases

upvoted 2 times

🗨️ **Maiana** 1 year, 5 months ago

ChatGPT: N/Y/Y

upvoted 3 times

🗨️ **ha1p** 1 year, 5 months ago

No, No, Yes

upvoted 2 times

🗨️ **XtraWest** 1 year, 5 months ago

as per my understanding: N | N | Y (Given answer is correct)

upvoted 2 times

  **samjoon** 1 year, 5 months ago

yny or nny?

upvoted 2 times

HOTSPOT -

Select the answer that correctly completes the sentence.

Hot Area:

### Answer Area

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

### Answer Area

Suggested Answer:

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

A graph database stores two types of information, nodes and edges. Edges specify relationships between nodes. Nodes and edges can have properties that provide information about that node or edge, similar to columns in a table. Edges can also have a direction indicating the nature of the relationship.

Reference:

<https://docs.microsoft.com/en-us/azure/architecture/guide/technology-choices/data-store-overview>

MEIQIN Highly Voted 1 year, 11 months ago

It means question 37 should be graph.  
upvoted 13 times

fusioner Highly Voted 2 years ago

Strange... Question 37 says something different...  
upvoted 11 times

okayhey Most Recent 6 months, 1 week ago

"A graph database stores two types of information, nodes and edges. Edges specify relationships between nodes. Nodes and edges can have properties that provide information about that node or edge, similar to columns in a table. Edges can also have a direction indicating the nature of the relationship."

<https://learn.microsoft.com/en-us/azure/architecture/guide/technology-choices/data-store-overview>

They're all RDS, but only graph supports analysis of relationship between entities

upvoted 2 times

AGTraining 6 months, 2 weeks ago

graph db  
upvoted 2 times

A7medM 7 months, 3 weeks ago

Duplicate with Q #20  
upvoted 1 times

MannyJZ 1 year, 1 month ago

On exam today 08/08/23  
upvoted 3 times

Gecig 1 year, 2 months ago

I don't see a question 37 which others are discussing about here. Is this visible only for contributor access??  
upvoted 7 times

gabo 7 months ago

I have contributor access and do not see Q 37, you are not alone.

upvoted 1 times

🗨️ **chiiiweiii** 1 year, 6 months ago

Correct Answer

upvoted 5 times

🗨️ **Joanna0** 1 year, 7 months ago

Correct answer.

Duplicated to question 37.

upvoted 2 times

🗨️ **Nurialzard** 1 year, 8 months ago

Same question as number 37

upvoted 2 times

🗨️ **Anil\_Mahanty** 1 year, 8 months ago

Relation between employee and entity, are two different things. entity relation belongs to graph and employee relation i.e. record relation

upvoted 2 times

🗨️ **SuleyDM** 1 year, 8 months ago

The question 37 has the wrong answer. Both in question 37 and 52 should be GRAPH

upvoted 6 times

🗨️ **edgarmassa** 1 year, 8 months ago

Galera, primeira é GRAFOS, e a resposta da outra é GRAFOS

upvoted 3 times

🗨️ **Dewonnaluv** 1 year, 9 months ago

Definitely, the earlier question 37 answer picked as Column family database was wrong.

upvoted 2 times

🗨️ **hungry85** 1 year, 9 months ago

GRAPH is the right answer

upvoted 3 times

🗨️ **Mannie\_001** 1 year, 10 months ago

ITS Graph Database Question number 37 has it wrong

upvoted 2 times

🗨️ **PatYeo** 1 year, 11 months ago

Graph database is correct.

upvoted 6 times

HOTSPOT -

Select the answer that correctly completes the sentence.

Hot Area:

### Answer Area

A key/value data store is optimized for

	▼
enforcing constraints.	
simple lookups.	
table joins.	

### Answer Area

**Suggested Answer:** A key/value data store is optimized for

	▼
enforcing constraints.	
simple lookups.	
table joins.	

A key/value store associates each data value with a unique key.

Key/value stores are highly optimized for applications performing simple lookups, but are less suitable if you need to query data across different key/value stores.

Key/value stores are also not optimized for querying by value.

A single key/value store can be extremely scalable, as the data store can easily distribute data across multiple nodes on separate machines.

Reference:

<https://docs.microsoft.com/en-us/azure/architecture/guide/technology-choices/data-store-overview>

- 👤 **nmbuff** Highly Voted 1 year, 11 months ago  
Simple Lookup is correct.  
upvoted 16 times
- 👤 **Albertology** Highly Voted 1 year, 11 months ago  
on exam 10 Oct 2022  
upvoted 5 times
- 👤 **heobanten04** Most Recent 5 months, 2 weeks ago  
I took the exam on 30.03.2024 and got 936; this question appeared, correct answer: Simple Lookups  
upvoted 1 times
- 👤 **AGTraining** 6 months, 2 weeks ago  
simple lookup  
upvoted 1 times
- 👤 **MannyJZ** 1 year, 1 month ago  
On exam today 08/08/23  
upvoted 1 times
- 👤 **Charley92** 1 year, 4 months ago  
on exam 05/20/2023  
upvoted 2 times
- 👤 **chiiiweiii** 1 year, 6 months ago  
Correct Answer  
upvoted 2 times
- 👤 **Nurialzard** 1 year, 8 months ago  
same question as number 20  
upvoted 2 times
- 👤 **PatYeo** 1 year, 11 months ago  
Simple Lookup is correct.

upvoted 3 times

 **fusioner** 2 years ago

Repeated Question.

upvoted 2 times

HOTSPOT -

Select the answer that correctly completes the sentence.

Hot Area:

### Answer Area

In an infrastructure as a service (IaaS) instance of  that hosts SQL Server

elastic pool
MySQL server
PostgreSQL server
virtual machine

### Answer Area

**Suggested Answer:** In an infrastructure as a service (IaaS) instance of  that hosts SQL Server

elastic pool
MySQL server
PostgreSQL server
virtual machine

SQL Server on Azure VM -

SQL Server on Azure VM falls into the industry category Infrastructure-as-a-Service (IaaS) and allows you to run SQL Server inside a fully managed virtual machine (VM) in Azure.

The most significant difference from SQL Database and SQL Managed Instance is that SQL Server on Azure Virtual Machines allows full control over the database engine. You can choose when to start maintenance/patching, change the recovery model to simple or bulk-logged, pause or start the service when needed, and you can fully customize the SQL Server database engine. With this additional control comes the added responsibility to manage the virtual machine.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/azure-sql-iaas-vs-paas-what-is-overview?view=azuresql>

-  **AndreLima** Highly Voted 1 year, 11 months ago  
Said VM, said IAAS.  
upvoted 24 times
-  **PatYeo** Highly Voted 1 year, 11 months ago  
d is correct. In IaaS, you need to manage the VM and OS.  
upvoted 9 times
-  **AGTraining** Most Recent 6 months, 2 weeks ago  
virtual m  
upvoted 1 times
-  **akhil5432** 1 year, 1 month ago  
IN IAAS WE MANAGE virtual machine  
upvoted 1 times
-  **chiiiweiii** 1 year, 6 months ago  
Correct Answer  
upvoted 3 times
-  **Dewonnaluv** 1 year, 9 months ago  
This also validates the community answer to an earlier question that assigned this as Software as a service(SaaS)  
upvoted 2 times
-  **Anumani** 1 year, 11 months ago  
d is correct  
upvoted 4 times
-  **fusioner** 2 years ago  
Correct Answer...  
upvoted 2 times

HOTSPOT -

Select the answer that correctly completes the sentence.

Hot Area:

### Answer Area

Varying fields for each entity in a JSON document is an example of

	▼
relational data.	
semi-structured data.	
structured data.	
unstructured data.	

### Answer Area

Suggested Answer:

Varying fields for each entity in a JSON document is an example of

	▼
relational data.	
semi-structured data.	
structured data.	
unstructured data.	

For some semi-structured data formats (e.g. JSON), data sets are frequently a simple concatenation of multiple documents.

Reference:

<https://docs.snowflake.com/en/user-guide/semistructured-considerations.html>

- 👤 **PatYeo** Highly Voted 1 year, 11 months ago  
 Semi-structured is correct. Semi-structured is supported via JSON, XML.  
 upvoted 25 times
- 👤 **Kermedi2018** Highly Voted 2 years ago  
 Correct  
 upvoted 7 times
- 👤 **okayhey** Most Recent 6 months, 1 week ago  
 Correct, varying fields means the data has a loose structure, meaning it's semi-structured  
 upvoted 1 times
- 👤 **AGTraining** 6 months, 2 weeks ago  
 Semi-structured  
 upvoted 1 times
- 👤 **trashbox** 10 months, 2 weeks ago  
 B. Semi-structured.
- Appeared on 2023-11-05.  
 upvoted 1 times
- 👤 **akhil5432** 1 year, 1 month ago  
 semi-structured data  
 upvoted 1 times
- 👤 **VNR25** 1 year, 2 months ago  
 This question was asked in a different way in the 07/07/23 exam and the answer is correct.  
 upvoted 1 times
- 👤 **chiiweiii** 1 year, 6 months ago  
 Correct Answer  
 upvoted 2 times
- 👤 **berend1** 1 year, 11 months ago  
 Correct  
 upvoted 5 times

HOTSPOT -

Select the answer that correctly completes the sentence.

Hot Area:

### Answer Area

A JSON document is an example of

graph data.
relational data.
semi-structured data.
unstructured data.

### Answer Area

Suggested Answer: A JSON document is an example of

graph data.
relational data.
semi-structured data.
unstructured data.

Semi-structured data -

Semi-structured data (e.g., JSON, CSV, XML) is the  $\lambda\epsilon\text{bridge}\lambda\epsilon$  between structured and unstructured data. It does not have a predefined data model and is more complex than structured data, yet easier to store than unstructured data.

Reference:

<https://www.ibm.com/cloud/blog/structured-vs-unstructured-data>

-  **PatYeo** Highly Voted 1 year, 11 months ago

Semi-structured in correct. JSON allows varying no. of key-values in each object.

upvoted 15 times
-  **AGTraining** Most Recent 6 months, 2 weeks ago

Semi-structured

upvoted 2 times
-  **VNR25** 1 year, 2 months ago

Q55 & 56 are more or less the same. This question appeared in the 07/07/23 exam and the answer is correct.

upvoted 3 times
-  **chiiiweiii** 1 year, 6 months ago

Correct Answer

upvoted 4 times
-  **Zaiky** 1 year, 10 months ago

If it's JASON, it's got to be semi structured!

upvoted 2 times
-  **TomGray1989** 1 year, 10 months ago

JSON files are examples of semi-structured data - answer is correct. If you didn't know this prior to reading this question, you need to get back to Microsoft learn or utilise the docs.

upvoted 1 times
-  **nmbuff** 1 year, 11 months ago

double question with a privious one

upvoted 1 times
-  **[Removed]** 2 years ago

Agree with the answer

upvoted 3 times

Which property of a transactional workload guarantees that each transaction is treated as a single unit that either succeeds completely or fails completely?

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

**Suggested Answer: A**

The purpose of making transactions atomic is to prevent different transactions from interfering with one another. This can only happen if more than one user process is trying to access the database at the same time, as when a server allows several clients to use it concurrently. The simplest way to enforce atomicity is for the DBMS to refuse to start any transaction until the previous one has committed.

Reference:

<https://www.sciencedirect.com/topics/computer-science/atomic-transaction>

Community vote distribution

A (100%)

 **PatYeo** Highly Voted 1 year, 11 months ago

A is correct.

<https://www.databricks.com/glossary/acid-transactions>

A.C.I.D. properties: Atomicity, Consistency, Isolation, and Durability

Atomicity - each statement in a transaction (to read, write, update or delete data) is treated as a single unit. Either the entire statement is executed, or none of it is executed. This property prevents data loss and corruption from occurring if, for example, if your streaming data source fails mid-stream.

upvoted 23 times

 **fusioner** Highly Voted 2 years ago

Correct Answer is A

upvoted 6 times

 **heobanten04** Most Recent 5 months, 2 weeks ago

I took the exam on 30.03.2024 and got 936; this question appeared, correct answer: Atomicity

upvoted 1 times

 **AGTraining** 6 months, 2 weeks ago

Selected Answer: A

option A

upvoted 1 times

 **trashbox** 10 months, 2 weeks ago

A.

Appeared on 2023-11-05.

upvoted 1 times

 **AZFabio** 1 year ago

Selected Answer: A

correct

upvoted 1 times

 **rohirag** 1 year, 2 months ago

Selected Answer: A

Atomicity

upvoted 1 times

🗨️ 👤 **VNR25** 1 year, 2 months ago

This question appeared in the 07/07/23 exam and the answer is correct.  
upvoted 1 times

🗨️ 👤 **mamahani** 1 year, 5 months ago

I took the exam 12.04.2023 and got 1000, 100% score; this question appeared, correct answer:  
atomicity  
upvoted 4 times

🗨️ 👤 **vizwiz** 1 year, 6 months ago

On exam March 9, 2023.  
upvoted 5 times

🗨️ 👤 **PhyoPluto** 1 year, 6 months ago

the exam on 07/03/2023  
upvoted 2 times

🗨️ 👤 **chiiiweiii** 1 year, 6 months ago

**Selected Answer: A**

Correct Answer  
upvoted 1 times

🗨️ 👤 **sreekan2** 1 year, 8 months ago

**Selected Answer: A**

A is correct.  
upvoted 1 times

🗨️ 👤 **Mannie\_001** 1 year, 10 months ago

Answer is correct  
upvoted 1 times

🗨️ 👤 **edurakhan** 1 year, 10 months ago

**Selected Answer: A**

Atomicity is the correct answer  
upvoted 2 times

🗨️ 👤 **JA2018** 1 year, 11 months ago

**Selected Answer: A**

all operations must be successfully completed for the transaction to complete.  
upvoted 3 times

🗨️ 👤 **raydel92** 1 year, 11 months ago

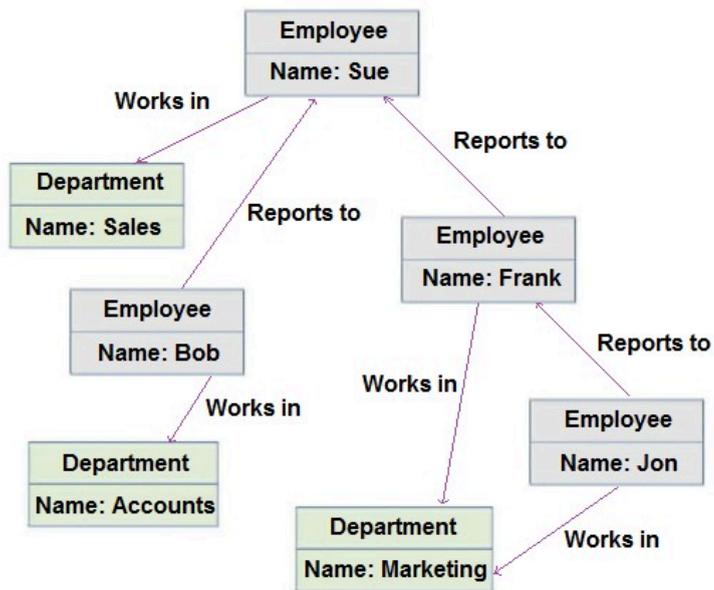
**Selected Answer: A**

Correct answer is A. atomicity

<https://learn.microsoft.com/en-us/training/modules/explore-core-data-concepts/5-transactional-data-processing?ns-enrollment-type=learningpath&ns-enrollment-id=learn.wvl.azure-data-fundamentals-explore-core-data-concepts>

upvoted 4 times

You have a data store that has the structure shown in the following exhibit.



Which type of data store is this?

- A. key/value
- B. object data
- C. graph
- D. time series

**Suggested Answer:** C

A graph database stores two types of information, nodes and edges. Edges specify relationships between nodes. Nodes and edges can have properties that provide information about that node or edge, similar to columns in a table. Edges can also have a direction indicating the nature of the relationship.

Reference:

<https://docs.microsoft.com/en-us/azure/architecture/guide/technology-choices/data-store-overview>

Community vote distribution

C (100%)

**Kermedi2018** Highly Voted 2 years ago

Selected Answer: C

Yes C is the correct answer  
upvoted 18 times

**fusioner** Highly Voted 2 years ago

Selected Answer: C

Yes, Graph  
upvoted 7 times

**AGTraining** Most Recent 6 months, 2 weeks ago

option C  
upvoted 1 times

**AZFabio** 1 year ago

Selected Answer: C

Graph is correct  
upvoted 1 times

**globy118** 1 year, 2 months ago

Was on exam 24/6/2023

upvoted 1 times

🗨️ **chiiiweiii** 1 year, 6 months ago

**Selected Answer: C**

Correct Answer

upvoted 3 times

🗨️ **Joanna0** 1 year, 7 months ago

**Selected Answer: C**

Graph.

<https://learn.microsoft.com/en-us/azure/architecture/guide/technology-choices/data-store-overview#graph-databases>

upvoted 4 times

🗨️ **sigh27** 1 year, 7 months ago

C is the correct answer

upvoted 1 times

🗨️ **Hari\_kg** 1 year, 7 months ago

**Selected Answer: C**

Graph database

upvoted 1 times

🗨️ **Yomzie** 7 months ago

Not "database" though. It is a Graph data model.

upvoted 1 times

🗨️ **edurakhan** 1 year, 10 months ago

**Selected Answer: C**

Graph database

upvoted 3 times

🗨️ **JA2018** 1 year, 11 months ago

**Selected Answer: C**

It's the only viable option.

upvoted 3 times

🗨️ **[Removed]** 2 years ago

**Selected Answer: C**

its graph

upvoted 4 times

🗨️ **Jiviify** 2 years ago

**Selected Answer: C**

Yes, its Graph

upvoted 4 times

Which type of database contains nodes and edges?

- A. graph
- B. key/value
- C. columnar
- D. time series

**Suggested Answer: A**

A graph database stores two types of information, nodes and edges. Edges specify relationships between nodes. Nodes and edges can have properties that provide information about that node or edge, similar to columns in a table. Edges can also have a direction indicating the nature of the relationship.

Reference:

<https://docs.microsoft.com/en-us/azure/architecture/guide/technology-choices/data-store-overview>

Community vote distribution

A (100%)

 **fusioner** Highly Voted 2 years ago

**Selected Answer: A**

Yes, Graph

upvoted 16 times

 **JA2018** Highly Voted 1 year, 11 months ago

**Selected Answer: A**

A graph database shows two types of information, nodes and edges. The edges specify the relationships between nodes. Both nodes and edges own properties that provide details about that node or edge. This is similar to columns in a table. The edges can also have a direction to indicate the relationship's nature.

upvoted 7 times

 **okayhey** Most Recent 6 months, 1 week ago

Si many graph questions, I hope I get them on my exam too :D

upvoted 1 times

 **AGTraining** 6 months, 2 weeks ago

option A

upvoted 1 times

 **trashbox** 10 months, 2 weeks ago

**Selected Answer: A**

Appeared on 2023-11-05.

upvoted 1 times

 **AZFabio** 1 year ago

**Selected Answer: A**

Graph is correct

upvoted 1 times

 **chiiiweiii** 1 year, 6 months ago

**Selected Answer: A**

Correct Answer

upvoted 2 times

 **Venky02** 1 year, 8 months ago

**Selected Answer: A**

A, GRAPH

upvoted 1 times

HOTSPOT -

Select the answer that correctly completes the sentence.

Hot Area:

### Answer Area

	▼	is responsible for identifying which business rules must be applied to the data of a company.
A data analyst		
A data engineer		
A data scientist		

### Answer Area

**Suggested Answer:**

	▼	is responsible for identifying which business rules must be applied to the data of a company.
A data analyst		
A data engineer		
A data scientist		

A data analyst's primary skill set revolves around data acquisition, handling, and processing.

Incorrect:

\* A data engineer requires an intermediate level understanding of programming to build thorough algorithms along with a mastery of statistics and math.

\* A data scientist needs to be a master of both worlds. Data, stats, and math along with in-depth programming knowledge for Machine Learning and Deep Learning.

Reference:

<https://www.edureka.co/blog/data-analyst-vs-data-engineer-vs-data-scientist/>

🗨️ **Massimo70** Highly Voted 1 year, 11 months ago

DATA ENGINEER: Data engineers manage infrastructure and processes for data integration across the organization, applying data cleaning routines, identifying data governance RULES, and implementing pipelines to transfer and transform data between systems.  
upvoted 56 times

🗨️ **\_Vijay\_** 1 month, 2 weeks ago

Data Engineer is not correct. This is the job of Data Analyst.  
upvoted 1 times

🗨️ **Kanatius** 1 year, 4 months ago

Governance rules and Business rules are 2 different things. GR are about who can access resource/datasets and who cannot, otherwise BR are about the solution rules. The correct answer is Data Analyst  
upvoted 11 times

🗨️ **p1zz4** 1 year, 8 months ago

Correct Answer : DATA ENGINEER  
upvoted 4 times

🗨️ **Pavi20210211** 1 year, 9 months ago

Correct Answer : DATA ENGINEER  
upvoted 6 times

🗨️ **AdminAy** Highly Voted 1 year, 11 months ago

Data Analyst analyzes numeric data and uses it to help companies make better decisions. Data Engineer involves in preparing data. They develop, constructs, tests & maintain complete architecture. A data scientist analyzes and interpret complex data.

Source: Google

upvoted 22 times

🗨️ **JoshPark21** Most Recent 1 month ago

I think, the answer is correct.

key job roles that deal with data in most organizations are:

- Database administrators manage databases, assigning permissions to users, storing backup copies of data and restore data in the event of a failure.

- Data engineers manage infrastructure and processes for data integration across the organization, applying data cleaning routines, identifying data governance rules, and implementing pipelines to transfer and transform data between systems.
- Data analysts explore and analyze data to create visualizations and charts that enable organizations to make informed decisions. processes raw data into relevant insights based on identified business requirements to deliver relevant insights.

upvoted 2 times

🗨️ 👤 **b9f0525** 3 months, 2 weeks ago

By ChatGPT4o:

A Data Analyst is responsible for identifying which business rules must be applied to the data of a company.

Explanation:

Data Engineer: Primarily focuses on designing, building, and maintaining the infrastructure and architecture for data generation, storage, and processing. They ensure that data pipelines are reliable and efficient but typically do not define business rules.

Data Analyst: Works closely with business stakeholders to understand their requirements and translate them into data-related tasks. They are responsible for analyzing data, identifying business rules, and ensuring that data is processed in a way that aligns with business needs.

Data Scientist: Uses statistical methods, algorithms, and machine learning to extract insights and build predictive models from data. While they may apply business rules in their analysis, defining these rules is generally not their primary responsibility.

upvoted 3 times

🗨️ 👤 **Koteyko** 6 months ago

In my practice Data Engineer (DE) never identifies rules, DE implements requirements provided by Business/Data Analyst (BA/DA). Identification of rules by communicating with Business should be a task of BA/DA.

upvoted 3 times

🗨️ 👤 **AGTraining** 6 months, 2 weeks ago

DATA ENGINEER

upvoted 1 times

🗨️ 👤 **Edgy\_San** 7 months, 1 week ago

I am divided. The question should either state governance rules or business requirements. Instead we are left to guess if requirement can be understood as business rules. Also governance rules usually is about standards that need to be met and compliance. To me at least.

upvoted 1 times

🗨️ 👤 **gcsilva83** 9 months, 2 weeks ago

A data analyst processes raw data into relevant insights based on identified business requirements to deliver relevant insights.

upvoted 1 times

🗨️ 👤 **AZFabio** 1 year ago

Very tricky as both roles can do it. But I choose Data Engineer as well, since he is the first in the chain on getting in touch with data and treating it (cleaning, dedupping, etc.).

upvoted 1 times

🗨️ 👤 **Sourabh808** 1 year, 1 month ago

yes correct

upvoted 2 times

🗨️ 👤 **srgopalkr** 1 year, 1 month ago

Tricky question as it depends on the contextual interpretation the term 'Business Rules'

Here's ChatGPT response "business rules relate to data quality in at least two fundamental ways. First, they can automate the decisions that the company makes in its day-to-day operations. Second, they can be used to audit data produced by existing processes for compliance with external regulation as well as internal business policies and goals." Both these sound to be the responsibility of the Data Analyst.

upvoted 2 times

🗨️ 👤 **VNR25** 1 year, 2 months ago

This question appeared in the 07/07/23 exam and the answer is correct.

upvoted 2 times

🗨️ 👤 **globy118** 1 year, 2 months ago

Was on exam 24/6/2023

upvoted 2 times

🗨️ 👤 **Nonin** 1 year, 2 months ago

and the correct answer?

upvoted 2 times

🗨️ 👤 **sanmighty** 1 year, 5 months ago

Data Engineer is correct

A data engineer collaborates with stakeholders to design and implement data-related workloads, including data ingestion pipelines, cleansing and transformation activities, and data stores for analytical workloads

<https://learn.microsoft.com/en-us/training/modules/explore-roles-responsibilities-world-of-data/2-explore-job-roles>

upvoted 5 times

🗨️ 👤 **XtraWest** 1 year, 5 months ago

A data analyst | as per chat gpt

upvoted 5 times

🗨️ 👤 **aspiringazureengineer** 1 year, 6 months ago

Data Engineers - manage infrastructure and processes for data integration across the organization, applying data cleaning routines, identifying data governance rules, and implementing pipelines to transfer and transform data between systems.

<https://learn.microsoft.com/en-us/training/modules/explore-roles-responsibilities-world-of-data/2-explore-job-roles>

upvoted 1 times

🗨️ 👤 **evanalytics** 1 year, 6 months ago

DATA ENGINEER: Los ingenieros de datos administran la infraestructura y los procesos para la integración de datos en una organización, aplican rutinas de limpieza de datos, IDENTIFICAN REGLAS de gobernanza de datos e implementan canalizaciones para transferir y transformar datos entre sistemas.

upvoted 1 times

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

**Suggested Answer:** C

Reference:

<https://docs.microsoft.com/en-us/sql/t-sql/queries/update-transact-sql?view=sql-server-ver15>

Community vote distribution

C (100%)

**vince60370** Highly Voted 3 years, 6 months ago

Update is the SIMPLEST way, but a MERGE can also do the update... I think the question should be a little bit more specific.  
upvoted 32 times

**DSilsbury** 2 years, 1 month ago

Merge requires 2 datasets, it can't be merge  
upvoted 7 times

**Armandoo** Highly Voted 3 years, 3 months ago

The answer is UPDATE since we want to change something within a column.  
upvoted 28 times

**Lipseal** 2 years, 4 months ago

I googled it:  
ALTER - add/delete/modify column  
UPDATE - update existing row  
INSERT - Insert new row  
upvoted 11 times

**AGTraining** Most Recent 6 months, 2 weeks ago

option C  
upvoted 1 times

**akhil5432** 1 year, 1 month ago

'Update' is correct option  
upvoted 1 times

**evanalytics** 1 year, 6 months ago

UPDATE  
upvoted 2 times

**chiiiweiii** 1 year, 6 months ago

**Selected Answer: C**  
Correct Answer  
upvoted 2 times

**JA2018** 1 year, 11 months ago

**Selected Answer: C**  
You need DML to update/ modify each row record in a DB table.

upvoted 2 times

🗨️ 👤 **Eltooth** 2 years, 2 months ago

**Selected Answer: C**

C is correct answer : Update

upvoted 4 times

🗨️ 👤 **alveswlm** 2 years, 2 months ago

**Selected Answer: C**

Update statement is used to change the value of a column.

upvoted 4 times

🗨️ 👤 **EXODUS33** 2 years, 4 months ago

**Selected Answer: C**

update is used to change an existing record value

upvoted 4 times

🗨️ 👤 **olusesan** 2 years, 4 months ago

Answer is C

upvoted 2 times

🗨️ 👤 **nhftloh** 2 years, 4 months ago

**Selected Answer: C**

Answer is c

upvoted 3 times

🗨️ 👤 **D\_Mo\_ZY** 2 years, 6 months ago

Update

Was in the exam (18/03/2022)

upvoted 5 times

🗨️ 👤 **Yvekke** 2 years, 8 months ago

update is ok

upvoted 2 times

🗨️ 👤 **jmay** 2 years, 9 months ago

even when using MERGE, you still need to include an UPDATE sub query

upvoted 5 times

🗨️ 👤 **Afamdi** 3 years ago

Update obviously!

upvoted 4 times

🗨️ 👤 **MFO\_FM** 3 years ago

Update

upvoted 3 times

Your company needs to implement a relational database in Azure. The solution must minimize ongoing maintenance. Which Azure service should you use?

- A. Azure HDInsight
- B. Azure SQL Database
- C. Azure Cosmos DB
- D. SQL Server on Azure Virtual Machines

**Suggested Answer:** B

Reference:

<https://azure.microsoft.com/en-us/services/sql-database/#features>

Community vote distribution

B (100%)

🗳️ 👤 **Armandoo** Highly Voted 3 years, 3 months ago

Correct Answer: B

Minimize maintenance: so it should a PaaS + Relational Database = Azure SQL Database  
upvoted 70 times

🗳️ 👤 **D\_Mo\_ZY** Highly Voted 2 years, 6 months ago

Correct Answer: B

Was in the exam (18/03/2022)  
upvoted 21 times

🗳️ 👤 **\_Vijay\_** Most Recent 1 month, 2 weeks ago

Selected Answer: B

Correct answer.  
upvoted 1 times

🗳️ 👤 **AGTraining** 6 months, 2 weeks ago

option B  
upvoted 1 times

🗳️ 👤 **AZFabio** 1 year ago

Selected Answer: B

Azure SQL DB is correct  
upvoted 1 times

🗳️ 👤 **VNR25** 1 year, 2 months ago

This question appeared in the 07/07/23 exam and the answer is correct.  
upvoted 1 times

🗳️ 👤 **HRKK** 1 year, 2 months ago

CORRECT B  
upvoted 1 times

🗳️ 👤 **Genius365** 1 year, 4 months ago

Selected Answer: B

Minimize maintenance: so it should a PaaS + Relational Database = Azure SQL Database  
upvoted 1 times

🗳️ 👤 **mamahani** 1 year, 5 months ago

I took the exam 12.04.2023 and got 1000, 100% score; this question appeared, correct answer:  
Azure SQL Database  
upvoted 7 times

🗳️ 👤 **chiiweiii** 1 year, 6 months ago

**Selected Answer: B**

Correct Answer  
upvoted 2 times

🗨️ 👤 **Joanna0** 1 year, 7 months ago

**Selected Answer: B**

Part of the Azure SQL family, Azure SQL Database is an always-up-to-date, fully managed relational database service built for the cloud.

<https://azure.microsoft.com/en-us/products/azure-sql/database/>  
upvoted 3 times

🗨️ 👤 **edurakhan** 1 year, 10 months ago

**Selected Answer: B**

B is the correct answer. Minimize maintenance is the key phrase  
upvoted 4 times

🗨️ 👤 **Albertology** 1 year, 11 months ago

on exam 10 Oct 2022

upvoted 3 times

🗨️ 👤 **JA2018** 1 year, 11 months ago

**Selected Answer: B**

PaaS = Minimized maintenance  
upvoted 4 times

🗨️ 👤 **Kermedi2018** 2 years ago

**Selected Answer: B**

Correct answer B  
upvoted 3 times

🗨️ 👤 **fusioner** 2 years ago

**Selected Answer: B**

Correct Answer  
upvoted 2 times

🗨️ 👤 **Eltooth** 2 years, 2 months ago

**Selected Answer: B**

B is correct answer : Azure SQL Database  
upvoted 2 times

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?

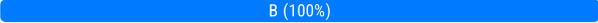
- A. Microsoft SQL Server Management Studio (SSMS)
- B. Azure Data Studio
- C. Azure CLI
- D. Azure PowerShell

**Suggested Answer: B**

Reference:

<https://www.mssqltips.com/sqlservertip/5997/create-sql-server-notebooks-in-azure-data-studio/>

Community vote distribution

 B (100%)

 **Fellipo** Highly Voted 3 years, 10 months ago

B its OK , <https://docs.microsoft.com/en-us/sql/azure-data-studio/notebooks/notebooks-sql-kernel?view=sql-server-ver15>  
upvoted 71 times

 **Ashodwbi** Highly Voted 3 years, 10 months ago

Answer is B  
upvoted 26 times

 **b9f0525** Most Recent 3 months, 2 weeks ago

Selected Answer: B  
Data Studio  
upvoted 1 times

 **AZFabio** 1 year ago

Selected Answer: B  
Data Studio  
upvoted 2 times

 **chiiiweiii** 1 year, 6 months ago

Selected Answer: B  
Correct Answer  
upvoted 4 times

 **Joanna0** 1 year, 7 months ago

Selected Answer: B  
Azure Data Studio

<https://learn.microsoft.com/en-us/sql/azure-data-studio/notebooks/notebooks-sql-kernel?view=sql-server-ver16>  
upvoted 3 times

 **Pinha** 1 year, 7 months ago

B. Azure Data Studio is a cross-platform database tool that allows you to embed documents and query results into a SQL notebook. It supports Azure SQL databases, SQL Server, and other databases. It also has built-in support for SQL Notebooks, which allows you to mix markdown, code, and results in a single notebook. This makes it a great tool for creating documentation and troubleshooting guides for administrators to use when working with Azure SQL databases.  
upvoted 14 times

 **iesRuben** 1 year, 9 months ago

Selected Answer: B  
B its the answer  
upvoted 3 times

 **JA2018** 1 year, 11 months ago

**Selected Answer: B**

Microsoft SQL Server Management Studio (SSMS) is an on-prem product. This exam focus more on SQL on Azure. My 2 cents' worth.  
upvoted 7 times

 **ColoradoChick** 1 year, 2 months ago

Just an FYI - SSMS is not just an on-prem product. SSMS can also connect to Azure (I've used it for the last 6 years on Azure projects). IaaS: SQL Server VM PaaS: Azure SQL DB. You can also connect to Synapse SQL Pools. SSMS has native integration with Data Studio. I still use SSMS for everything because I do DBA/arch/dev work across on-prem & cloud. But Data Studio is cross-platform & has real coding languages.

This is a nice Azure Data Studio vs SSMS comparison (and which one is better for what type of task):

<https://mindmajix.com/azure-data-studio-vs-ssms>

... but, yes, the answer to this question itself is Data Studio.

Just wanted people to know that they can still use SSMS in Azure. The tool is exactly the same regardless of on-prem or cloud.

upvoted 3 times

 **Bambi0074** 1 year, 11 months ago

It is still B but if you use the current SSMS you'll get the Az Data Studio automatically ;-)

<https://learn.microsoft.com/en-us/sql/azure-data-studio/download-azure-data-studio?view=sql-server-ver16#windows-installation>

upvoted 5 times

 **Eltooth** 2 years, 2 months ago

**Selected Answer: B**

B is correct answer : Data studio

upvoted 1 times

 **cormorant** 2 years, 5 months ago

use data studio is you need top embed documents and query results into an sql notebook

upvoted 5 times

 **RichMelbourne** 2 years, 5 months ago

Shown on test 15/04/2022

upvoted 3 times

 **Ashutosh344** 2 years, 6 months ago

B is correct

upvoted 1 times

 **theLazyGuy** 2 years, 10 months ago

<https://docs.microsoft.com/en-us/sql/azure-data-studio/notebooks/notebooks-guidance?view=sql-server-ver15>

upvoted 3 times

 **xiban** 2 years, 11 months ago

Databricks

upvoted 2 times

 **cadio30** 3 years ago

Azure Studio application is correct for this requirement

upvoted 2 times

DRAG DROP -

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.

Select and Place:

### 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

### Suggested Answer:

#### 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

Reference:

[https://en.wikipedia.org/wiki/Table\\_\(database\)](https://en.wikipedia.org/wiki/Table_(database))

[https://en.wikipedia.org/wiki/View\\_\(SQL\)](https://en.wikipedia.org/wiki/View_(SQL))

[https://en.wikipedia.org/wiki/Database\\_index](https://en.wikipedia.org/wiki/Database_index)

 **S3** Highly Voted 3 years, 9 months ago

correct answer!

upvoted 73 times

 **mrjeet** Highly Voted 2 years, 8 months ago

Yep, all correct!

Table

View

Index

upvoted 15 times

 **QeUtPat** Most Recent 5 months, 2 weeks ago

T, V, I is the correct answer

upvoted 1 times

 **AGTraining** 6 months, 2 weeks ago

Table

View

Index

upvoted 1 times

 **gabo** 7 months ago

It's not needed to comment when you feel the answer provided is correct, that way we avoid unnecessarily reading comments.

upvoted 5 times

 **Someone728** 3 weeks, 5 days ago

I would say it is better, as we could have a question which might have a wrong answer provided, no harm in it tbh

upvoted 1 times

 **gustavobabolim** 1 year ago

correct

upvoted 1 times

🗨️ **globy118** 1 year, 2 months ago

Was on exam 24/6/2023

upvoted 1 times

🗨️ **chiiiweiii** 1 year, 6 months ago

Correct Answer

upvoted 2 times

🗨️ **nmbuff** 1 year, 11 months ago

Yep, all correct!

Table

View

Index

upvoted 2 times

🗨️ **fusioner** 2 years ago

Table, View, Index

upvoted 2 times

🗨️ **noussa** 2 years, 1 month ago

Table, View, Index

upvoted 2 times

🗨️ **ria97** 2 years, 1 month ago

Was in exam 11/08. Correct answer

upvoted 4 times

🗨️ **Breeze21** 2 years, 2 months ago

Exam: 19.07.22

upvoted 3 times

🗨️ **Eltooth** 2 years, 2 months ago

Table

View

Index

upvoted 3 times

🗨️ **mre** 2 years, 5 months ago

Correct

upvoted 4 times

🗨️ **Ganaps** 2 years, 5 months ago

CORRECT

upvoted 3 times

🗨️ **SiDoCiOuS** 2 years, 5 months ago

Appeared on the exam on 03/22/2022

Answer is correct!

upvoted 4 times

You have an e-commerce application that reads and writes data to an Azure SQL database.  
Which type of processing does the application use?

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

**Suggested Answer: D**

OLTP is designed to serve as a persistent data store for business or front-end applications. OLTP administers day to day transaction of an organization.

Reference:

<https://sqlwizard.blog/2020/03/15/sql-server-oltp-vs-olap/>

Community vote distribution

D (100%)

 **solozano0725** Highly Voted 3 years, 5 months ago

Correct! It's a transactional process.  
upvoted 45 times

 **mrjeet** Highly Voted 2 years, 8 months ago

Correct answer - OLTP

OLTP or Online Transaction Processing is a type of data processing that consists of executing a number of transactions occurring concurrently  
upvoted 11 times

 **QeUtPat** Most Recent 5 months, 2 weeks ago

OLTP is correct answer.  
upvoted 1 times

 **AGTraining** 6 months, 2 weeks ago

option D  
upvoted 1 times

 **gustavobolim** 1 year ago

resposta correta  
upvoted 2 times

 **akhil5432** 1 year, 1 month ago

Selected Answer: D  
online transaction processing  
upvoted 1 times

 **mamahani** 1 year, 5 months ago

I took the exam 12.04.2023; this question appeared,correct answer  
Online Transation Processing OLTP  
upvoted 5 times

 **chiiiweiii** 1 year, 6 months ago

Selected Answer: D  
Correct Answer  
upvoted 1 times

 **raydel92** 1 year, 11 months ago

Selected Answer: D  
Online Transaction Processing (OLTP)  
<https://learn.microsoft.com/en-us/training/modules/explore-core-data-concepts/5-transactional-data-processing?ns-enrollment-type=learningpath&ns-enrollment-id=learn.wwl.azure-data-fundamentals-explore-core-data-concepts>

upvoted 6 times

🗨️ 👤 **Eltooth** 2 years, 2 months ago

**Selected Answer: D**

D is correct answer.

upvoted 4 times

🗨️ 👤 **dk87** 2 years, 4 months ago

Exam: 05/18/2022

upvoted 4 times

🗨️ 👤 **heenak55** 2 years, 4 months ago

In exam 21 April 2022

upvoted 5 times

🗨️ 👤 **ZeroMos** 2 years, 5 months ago

**Selected Answer: D**

Correct

upvoted 4 times

🗨️ 👤 **CJSingh** 2 years, 5 months ago

Correct

upvoted 2 times

🗨️ 👤 **Roshan142** 2 years, 6 months ago

**Selected Answer: D**

CORRECT !!

upvoted 3 times

🗨️ 👤 **CrazyAunt** 2 years, 6 months ago

**Selected Answer: D**

Processing transactions for an ecommerce site so OLTP.

upvoted 5 times

🗨️ 👤 **Koph** 2 years, 9 months ago

OLTP is correct

upvoted 2 times

When can you use an Azure Resource Manager template?

- A. to automate the creation of an interdependent group of Azure resources in a repeatable way
- B. to apply Azure policies for multi-tenant deployments
- C. to provision Azure subscriptions
- D. to control which services and feature administrators and developers can deploy from the Azure portal

**Suggested Answer: A**

You can automate deployments and use the practice of infrastructure as code. In code, you define the infrastructure that needs to be deployed

To implement infrastructure as code for your Azure solutions, use Azure Resource Manager templates (ARM templates). The template is a JavaScript Object

Notation (JSON) file that defines the infrastructure and configuration for your project. The template uses declarative syntax, which lets you state what you intend to deploy without having to write the sequence of programming commands to create it. In the template, you specify the resources to deploy and the properties for those resources.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-resource-manager/templates/overview>

Community vote distribution

A (100%)

🗳️ 👤 **AlfuryDB** Highly Voted 👍 3 years, 6 months ago

This one is correct as well. ARM allows you to automate the creation of resources by specifying in a JSON document the resources you want to generate. The process can be replicated (if needed) to deploy resources to a different region.

upvoted 57 times

🗳️ 👤 **choyy16** Highly Voted 👍 3 years, 2 months ago

Correct! Similar question in AZ and AI 900 as well

upvoted 23 times

🗳️ 👤 **AGTraining** Most Recent 🕒 6 months, 2 weeks ago

option A

upvoted 1 times

🗳️ 👤 **PDR** 1 year, 2 months ago

I would of course choose A in the exam , but literally speaking all of those answers are correct, as for example you could parametrise the a template to deploy to multi tenant, azure policy can be deployed using ARM templates, subscriptions can be created with ARM.

upvoted 2 times

🗳️ 👤 **Kets\_1406** 1 year, 3 months ago

Selected Answer: A

A is correct

upvoted 1 times

🗳️ 👤 **chiiiweiii** 1 year, 6 months ago

Selected Answer: A

Correct Answer

upvoted 3 times

🗳️ 👤 **iesRuben** 1 year, 9 months ago

Selected Answer: A

A is correct

upvoted 1 times

🗳️ 👤 **Eltooth** 2 years, 2 months ago

Selected Answer: A

A is correct answer however this is outdated.

upvoted 6 times

🗨️ 👤 **bfan95** 2 years, 4 months ago

Selected Answer: A

A is correct

upvoted 4 times

🗨️ 👤 **mre** 2 years, 5 months ago

A is correct

upvoted 2 times

🗨️ 👤 **weri** 2 years, 5 months ago

Isn't D the correct answer?

upvoted 1 times

🗨️ 👤 **pads98** 2 years, 4 months ago

D is RBAC.

upvoted 2 times

🗨️ 👤 **mrjeet** 2 years, 8 months ago

A is correct - ARM templates allow you to create and deploy an entire Azure infrastructure declaratively

upvoted 4 times

🗨️ 👤 **Koph** 2 years, 9 months ago

A is correct

upvoted 2 times

🗨️ 👤 **PanP** 2 years, 10 months ago

Correct, but don't know why related to DP900

upvoted 11 times

🗨️ 👤 **Forspotboy** 3 years, 3 months ago

Correct!

upvoted 5 times

🗨️ 👤 **ANath** 3 years, 5 months ago

Correct

upvoted 3 times

You have an Azure SQL database that you access directly from the Internet.  
You recently changed the public IP address of your computer.  
After changing the IP address, you can no longer access the database. You can connect to other resources in Azure.  
What is a possible cause of the issue?

- A. role-based access control (RBAC)
- B. Dynamic Host Configuration Protocol (DHCP)
- C. Domain Name Service (DNS)
- D. a database-level firewall

**Suggested Answer: D**

The Azure SQL Database firewall lets you decide which IP addresses may or may not have access to either your Azure SQL Server or your Azure SQL database.

When creating an Azure SQL Database, the firewall needs to be configured before anyone will be able to access the database. By default, no external access to your SQL Database will be allowed until you explicitly assign permission by creating a firewall rule.

Reference:

<https://www.sqlshack.com/configuring-the-azure-sql-database-firewall/>

Community vote distribution

D (100%)

🗳️ **svetlanam** Highly Voted 2 years, 4 months ago

Correct. Answer is D  
upvoted 24 times

🗳️ **AlbertKwan** Highly Voted 1 year, 10 months ago

Selected Answer: D

D is correct. But this question shouldn't be in scope of this exam!  
upvoted 14 times

🗳️ **QeUtPat** Most Recent 5 months, 2 weeks ago

D is the correct answer  
upvoted 1 times

🗳️ **AGTraining** 6 months, 2 weeks ago

option D  
upvoted 1 times

🗳️ **chiiiweiii** 1 year, 6 months ago

Selected Answer: D

Correct Answer  
upvoted 2 times

🗳️ **nmbuff** 1 year, 11 months ago

Selected Answer: D

Yes. answer is D.  
Not A because  
Role-based access control (RBAC) is a method of restricting network access based on the roles of individual users within an enterprise.  
upvoted 7 times

🗳️ **PatYeo** 1 year, 11 months ago

The new IP address needs to be allowed in the Firewall rules. Correct Answer is D.  
upvoted 6 times

🗳️ **Anumani** 1 year, 11 months ago

Selected Answer: D

d is absolutely  
upvoted 2 times

🗨️ 👤 **noussa** 2 years, 1 month ago

Correct

upvoted 1 times

🗨️ 👤 **Eltooth** 2 years, 2 months ago

**Selected Answer: D**

D is correct answer. Db firewall

upvoted 2 times

🗨️ 👤 **simplementeluca** 2 years, 2 months ago

DHCP no?

upvoted 1 times

🗨️ 👤 **Tweet12** 2 years, 3 months ago

**Selected Answer: D**

D is the answer.

upvoted 1 times

🗨️ 👤 **bfan95** 2 years, 4 months ago

**Selected Answer: D**

Answer: D

upvoted 3 times

DRAG DROP -

Match the tools to the appropriate descriptions.

To answer, drag the appropriate tool from the column on the left to its description on the right. Each tool may be used once, more than once, or not at all.

Select and Place:

#### Tools

Azure Data Studio

Microsoft SQL Server Data Tools (SSDT)

Microsoft SQL Server Management Studio (SSMS)

#### Answer Area

Tool

A graphical tool for managing SQL Server or Azure SQL databases that supports access, configuration, management, and administration tasks.

Tool

A lightweight editor that can run on-demand SQL queries and view and save results as text, JSON, or Microsoft Excel files.

Tool

A development tool for building Azure SQL databases, Microsoft SQL Server relational databases, SQL Server Analysis Services (SSAS) data models, SQL Server Integration Services (SSIS) packages, and SQL Server Reporting Services (SSRS) reports.

#### Suggested Answer:

#### Tools

#### Answer Area

Microsoft SQL Server Management Studio (SSMS)

A graphical tool for managing SQL Server or Azure SQL databases that supports access, configuration, management, and administration tasks.

Azure Data Studio

A lightweight editor that can run on-demand SQL queries and view and save results as text, JSON, or Microsoft Excel files.

Microsoft SQL Server Data Tools (SSDT)

A development tool for building Azure SQL databases, Microsoft SQL Server relational databases, SQL Server Analysis Services (SSAS) data models, SQL Server Integration Services (SSIS) packages, and SQL Server Reporting Services (SSRS) reports.

Box 1: Microsoft SQL Server Management Studio (SSMS)

SQL Server Management Studio (SSMS) is an integrated environment for managing any SQL infrastructure, from SQL Server to Azure SQL Database.

Box 2: Azure Data Studio -

Azure Data Studio offers a modern, keyboard-focused SQL coding experience that makes your everyday tasks easier with built-in features, such as multiple tab windows, a rich SQL editor, IntelliSense, keyword completion, code snippets, code navigation, and source control integration (Git). Run on-demand SQL queries, view and save results as text, JSON, or Excel. Edit data, organize your favorite database connections, and browse database objects in a familiar object browsing experience.

Box 3: Microsoft SQL Server Data Tools (SSDT)

SQL Server Data Tools (SSDT) is a modern development tool for building SQL Server relational databases, databases in Azure SQL, Analysis Services (AS) data models, Integration Services (IS) packages, and Reporting Services (RS) reports. With SSDT, you can design and deploy any SQL Server content type with the same ease as you would develop an application in Visual Studio.

Reference:

<https://docs.microsoft.com/en-us/sql/ssms/download-sql-server-management-studio-ssms>

<https://code.visualstudio.com/docs/supporting/FAQ> <https://docs.microsoft.com/en-us/sql/azure-data-studio/what-is-azure-data-studio>

<https://docs.microsoft.com/en-us/sql/ssdt/download-sql-server-data-tools-ssdt>

 **fusioner** Highly Voted 2 years ago

SSMS, Azure Data Studio, SSDT

upvoted 24 times

 **PatYeo** Highly Voted 1 year, 11 months ago

SSDT for database projects is normally what people talk about when they say "why should I use ssdt over ssms", what I normally say is that ssms is an administration tool, and ssdt is the developer tool. If you have a simple database and it is not very complex then ssdt is probably an overhead you don't need.

<https://stackoverflow.com/questions/36370135/i-do-not-understand-what-ssdt-do-as-opposed-to-or-not-ssms-and-sql-server-obj#:~:text=SSDT%20for%20database%20projects%20is,overhead%20you%20don't%20need.>

upvoted 13 times

 **PI\_Team** Most Recent 4 months ago

You can see this table, also you know that you have access to this table during the exam

<https://learn.microsoft.com/en-us/sql/tools/overview-sql-tools?view=sql-server-ver16>

upvoted 3 times

🗨️ 👤 **AGTraining** 6 months, 2 weeks ago

SSMS,

Azure Data Studio,

SSDT

upvoted 1 times

🗨️ 👤 **HRKK** 1 year, 2 months ago

CORRECT ANSWER

upvoted 1 times

🗨️ 👤 **chiiweiii** 1 year, 6 months ago

Correct Answer

upvoted 3 times

🗨️ 👤 **PRALLAB** 1 year, 10 months ago

<https://learn.microsoft.com/en-us/sql/azure-data-studio/what-is-azure-data-studio?view=sql-server-ver15>.

Azure Data Studio offers a modern, keyboard-focused SQL coding experience that makes your everyday tasks easier with built-in features, such as multiple tab windows, a rich SQL editor, IntelliSense, keyword completion, code snippets, code navigation, and source control integration (Git). Run on-demand SQL queries, and view and save results as text, JSON, or Excel. Edit data, organize your favourite database connections and browse database objects in a familiar object browsing experience.

upvoted 4 times

HOTSPOT -

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE: Each correct selection is worth one point.

Hot Area:

### Answer Area

Statements	Yes	No
Relational database tables contain columns and rows	<input type="radio"/>	<input type="radio"/>
Indexes in a relational database describe the data types in a table	<input type="radio"/>	<input type="radio"/>
A database view is a virtual table whose content is defined by a query	<input type="radio"/>	<input type="radio"/>

Suggested Answer:

### Answer Area

Statements	Yes	No
Relational database tables contain columns and rows	<input checked="" type="radio"/>	<input type="radio"/>
Indexes in a relational database describe the data types in a table	<input type="radio"/>	<input checked="" type="radio"/>
A database view is a virtual table whose content is defined by a query	<input checked="" type="radio"/>	<input type="radio"/>

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.

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/tables/tables> <https://docs.microsoft.com/en-us/sql/relational-databases/indexes/clustered-and-nonclustered-indexes-described> <https://docs.microsoft.com/en-us/sql/relational-databases/views/views?view=sql-server-ver15>

 **kaleidoscope** Highly Voted 3 years, 3 months ago

Correct

upvoted 76 times

 **cianryan** Highly Voted 2 years, 10 months ago

YES, NO, YES

upvoted 34 times

 **heobanten04** Most Recent 5 months, 2 weeks ago

I took the exam on 30.03.2024 and got 936; this question appeared, correct answer: Yes, No, yes

upvoted 1 times

 **AGTraining** 6 months, 2 weeks ago

YES, NO, YES

upvoted 1 times

🗨️ **[Removed]** 1 year, 4 months ago

Yes, No. Yes! yayyyyyyyyy

upvoted 4 times

🗨️ **mamahani** 1 year, 5 months ago

I took the exam 12.04.2023 and got 1000, 100% score; this question appeared, correct answer:

yes / no / yes

upvoted 6 times

🗨️ **chiiweiii** 1 year, 6 months ago

YNY are Correct Answer

upvoted 2 times

🗨️ **Motivator** 2 years ago

Agreed

YES, NO, YES

upvoted 5 times

🗨️ **noussa** 2 years, 1 month ago

Correct

upvoted 4 times

🗨️ **dk87** 2 years, 4 months ago

Exam: 05/18/2022

upvoted 4 times

🗨️ **holyyyyyyyy** 2 years, 7 months ago

YNS is the correct ans

upvoted 4 times

🗨️ **G\_unit\_19** 2 years, 7 months ago

YES , NO , YES

upvoted 3 times

🗨️ **AzureDP900** 2 years, 11 months ago

Correct answer

upvoted 2 times

🗨️ **MFO\_FM** 3 years ago

Correct

upvoted 3 times

🗨️ **bettz\_druid** 3 years ago

agreed with the answers

upvoted 3 times

🗨️ **kashi1983** 3 years, 1 month ago

Answer is correct

upvoted 3 times

🗨️ **Durgesh09** 3 years, 1 month ago

correct

upvoted 2 times

Which command-line tool can you use to query Azure SQL databases?

- A. sqlcmd
- B. bcp
- C. azdata
- D. Azure CLI

**Suggested Answer: A**

The sqlcmd utility lets you enter Transact-SQL statements, system procedures, and script files at the command prompt.

Incorrect Answers:

B: The bulk copy program utility (bcp) bulk copies data between an instance of Microsoft SQL Server and a data file in a user-specified format.

D: The Azure CLI is the defacto tool for cross-platform and command-line tools for building and managing Azure resources.

Reference:

<https://docs.microsoft.com/en-us/sql/tools/overview-sql-tools?view=sql-server-ver15>

Community vote distribution

A (100%)

  **Sc2** Highly Voted 4 years ago

I think 90% the answer is A

<https://docs.microsoft.com/en-us/sql/ssms/scripting/sqlcmd-use-the-utility?view=sql-server-ver15>

Az cli is perfect to create and manage the DB but to query sqlcmd is the tool

upvoted 111 times

  **sreekarv** 3 years, 9 months ago

100% the answer is A

upvoted 32 times

  **Haidardba** 3 years, 9 months ago

Yes, Use the [sqlcmd][sqlcmd] command-line utility to connect to and query a SQL pool. <https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-get-started-connect-sqlcmd>

upvoted 7 times

  **ExamPwnr** Highly Voted 4 years ago

Correct answer: A

Azure CLI page: <https://docs.microsoft.com/en-us/azure/azure-sql/database/az-cli-script-samples-content-guide?tabs=single-database>

"You can configure Azure SQL Database and SQL Managed Instance by using the Azure CLI."

You cannot execute queries though.

sqlcmd: <https://docs.microsoft.com/en-us/sql/tools/sqlcmd-utility?view=sql-server-ver15>

It is clearly mentioned that Azure SQL is supported, and Query examples are provided in the link above.

upvoted 35 times

  **AGTraining** Most Recent 6 months, 2 weeks ago

option A

upvoted 1 times

  **indiaraj460** 9 months ago

F off

upvoted 1 times

  **akhil5432** 1 year, 1 month ago

Selected Answer: A

option A sqlcmd

upvoted 1 times

🗨️ **chiiweiii** 1 year, 6 months ago

Selected Answer: A

Correct Answer  
upvoted 4 times

🗨️ **iesRuben** 1 year, 9 months ago

Selected Answer: A

sqlcmd is correct  
upvoted 5 times

🗨️ **AlbertKwan** 1 year, 10 months ago

Selected Answer: A

A is good.

The sqlcmd utility lets you enter Transact-SQL statements, system procedures, and script files through a variety of available modes:

At the command prompt.

In Query Editor in SQLCMD mode.

In a Windows script file.

In an operating system (Cmd.exe) job step of a SQL Server Agent job.

upvoted 2 times

🗨️ **Vergee** 2 years, 2 months ago

sqlcmd is the correct answer!

upvoted 3 times

🗨️ **simplex06** 2 years, 3 months ago

"Azure Data CLI (azdata) is a command-line utility written in Python to bootstrap and manage the data services via REST APIs."

Reference: <https://docs.microsoft.com/en-us/sql/azdata/install/deploy-install-azdata?view=sql-server-ver16>

upvoted 1 times

🗨️ **bfan95** 2 years, 4 months ago

Selected Answer: A

sqlcmd

upvoted 2 times

🗨️ **willyto** 2 years, 8 months ago

Selected Answer: A

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-get-started-connect-sqlcmd>

upvoted 2 times

🗨️ **mudassirmsk27** 2 years, 8 months ago

A is the answer

upvoted 2 times

🗨️ **Haidardba** 2 years, 9 months ago

Selected Answer: A

The sqlcmd utility lets you enter Transact-SQL statements, system procedures, and script files through a variety of available modes:

At the command prompt.

In Query Editor in SQLCMD mode.

In a Windows script file.

In an operating system (Cmd.exe) job step of a SQL Server Agent job.

<https://docs.microsoft.com/en-us/sql/tools/sqlcmd-utility?view=sql-server-ver15>

upvoted 2 times

🗨️ **Haidardba** 2 years, 9 months ago

The Correct Answer is A.

upvoted 1 times

🗨️ **cadiao30** 3 years ago

A is the correct answer as it is use to "query" and not to perform administrative or provisioning of the service

upvoted 3 times

🗨️ **Feliasg9** 3 years ago

It's to QUERY, not DEPLOY, Answer is correct

upvoted 3 times

HOTSPOT -

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE: Each correct selection is worth one point.

Hot Area:

## Answer Area

Statements	Yes	No
Azure SQL Database includes a managed backup service.	<input type="radio"/>	<input type="radio"/>
Azure SQL Database has built-in high availability.	<input type="radio"/>	<input type="radio"/>
Azure SQL Database can use Azure Defender.	<input type="radio"/>	<input type="radio"/>

Suggested Answer:

## Answer Area

Statements	Yes	No
Azure SQL Database includes a managed backup service.	<input checked="" type="radio"/>	<input type="radio"/>
Azure SQL Database has built-in high availability.	<input checked="" type="radio"/>	<input type="radio"/>
Azure SQL Database can use Azure Defender.	<input checked="" type="radio"/>	<input type="radio"/>

Box 1: Yes -

Box 2: Yes -

Box 3: Yes -

Azure Defender provides security alerts and advanced threat protection for virtual machines, SQL databases, containers, web applications, your network, and more.

Azure Defender provides security alerts and advanced threat protection for virtual machines, SQL databases, containers, web applications, your network, and more.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/sql-database-paas-overview> <https://azure.microsoft.com/en-us/blog/announcing-sql-atp-and-sql-vulnerability-assessment-general-availability/> <https://docs.microsoft.com/en-us/azure/security-center/azure-defender>

 **AzureDP900** Highly Voted 2 years, 11 months ago

Yes, Yes. Yes

upvoted 53 times

 **Ihermos** 2 years, 11 months ago

It sounds like an ad to Azure SQL database

upvoted 73 times

  **girideshi** 2 years, 10 months ago

To be honest whole Azure certifications are marketing gimmick.

upvoted 74 times

  **impratik7** 2 years, 8 months ago

Couldn't agree more!

upvoted 11 times

  **sparkchu** 2 years, 6 months ago

this answer thread is a bit different that others...lol

upvoted 9 times

  **TouhiD** 2 years, 5 months ago

I agree with all 900 level courses, that is their purpose. but upper level courses like AZ104 and AZ500 are really good and if you have ESI account with your company, their labs are top notch !

upvoted 8 times

  **ColoradoChick** 1 year, 2 months ago

I agree TouhiD ... 900 Fundamentals certs are just that ... fundamentals. 900-level certs are like a first semester freshman year course - very basic, non-tech intro certs so you can learn the resources/tools.

Anyone who says all Microsoft certs are "a marketing gimmick" hasn't taken a technically-based MS cert (I have 4 on-prem certs). So unless the Azure certs are somehow easier than the on-prem certs were, the certs after 900 won't be a piece of cake like this.

I decided to do the 3 Fundamentals because I got free MS exam vouchers & it made me to learn basics of resources that will be hit technically in AZ-305 and DP-203.

upvoted 2 times

  **G\_unit\_19**  2 years, 7 months ago

Yes to all ... "come use this service, it has it all".

upvoted 7 times

  **AGTraining**  6 months, 2 weeks ago

Yes, Yes. Yes

upvoted 1 times

  **chiiiweiii** 1 year, 6 months ago

Correct Answer

upvoted 2 times

  **raydel92** 1 year, 11 months ago

Yes, yes, yes

<https://learn.microsoft.com/en-us/training/modules/explore-provision-deploy-relational-database-offerings-azure/2-azure-sql?ns-enrollment-type=learningpath&ns-enrollment-id=learn.wwl.azure-data-fundamentals-explore-relational-data>

upvoted 7 times

  **nikhilambati09** 2 years, 2 months ago

Yes, yes, yes

upvoted 2 times

  **bikewun** 2 years, 5 months ago

correct answer

upvoted 2 times

  **Charly6** 2 years, 5 months ago

Just say Yes!

upvoted 5 times

  **Koph** 2 years, 9 months ago

Answer is correct

upvoted 1 times

HOTSPOT -

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE: Each correct selection is worth one point.

Hot Area:

### Answer Area

Statements	Yes	No
You can use Azure Data Studio to query a Microsoft SQL Server big data cluster.	<input type="radio"/>	<input type="radio"/>
You can use Microsoft SQL Server Management Studio (SSMS) to query an Azure Synapse Analytics data warehouse.	<input type="radio"/>	<input type="radio"/>
You can use MySQL Workbench to query Azure Database for MariaDB databases.	<input type="radio"/>	<input type="radio"/>

### Answer Area

Statements	Yes	No
You can use Azure Data Studio to query a Microsoft SQL Server big data cluster.	<input checked="" type="radio"/>	<input type="radio"/>
You can use Microsoft SQL Server Management Studio (SSMS) to query an Azure Synapse Analytics data warehouse.	<input checked="" type="radio"/>	<input type="radio"/>
You can use MySQL Workbench to query Azure Database for MariaDB databases.	<input checked="" type="radio"/>	<input type="radio"/>

Suggested Answer:

Reference:

<https://docs.microsoft.com/en-us/sql/big-data-cluster/connect-to-big-data-cluster?view=sql-server-ver15> <https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-query-ssms> <https://docs.microsoft.com/en-us/azure/mariadb/connect-workbench>

learner301 Highly Voted 4 years ago

A) Yes from <https://docs.microsoft.com/en-us/sql/azure-data-studio/what-is?view=sql-server-ver15>

Use Azure Data Studio if you: Are connecting to a SQL Server 2019 big data cluster

B) Yes from <https://docs.microsoft.com/en-us/learn/modules/query-azure-sql-data-warehouse/4-query-dw-using-ssms>

C) Yes from <https://docs.microsoft.com/en-us/azure/mariadb/connect-workbench>

upvoted 153 times

niw66966 2 years, 9 months ago

Beautiful backup to each answer.

upvoted 17 times

Iroshan4 3 years, 4 months ago

Agreed

upvoted 5 times

LarsLarsgren Highly Voted 2 years, 1 month ago

Hey guys,

passed the exam today.

41 questions, 45 min, 700 points for passing.

Many questions were same like here in forum. Tq guys, you are gorgeous!

upvoted 28 times

AGTraining Most Recent 6 months, 2 weeks ago

Yes, Yes. Yes

upvoted 1 times

🗨️ 👤 **chiiweiii** 1 year, 6 months ago

Correct Answer

upvoted 3 times

🗨️ 👤 **LeonardoTran** 2 years ago

YES, YES, YES

upvoted 5 times

🗨️ 👤 **Deliboi** 2 years, 2 months ago

lol another marketing scheme for this..

upvoted 12 times

🗨️ 👤 **Armandoo** 3 years, 3 months ago

YES, YES, YES

upvoted 9 times

🗨️ 👤 **Jenny** 3 years, 7 months ago

Link to Quickstart guide for Azure Database for Maria DB: Use MySQL Workbench to connect to query data <https://docs.microsoft.com/en-us/azure/mariadb/quickstart-guide>  
workbench#:~:text=To%20connect%20to%20an%20Azure%20Database%20for%20MariaDB,all%20parameters%20are%20configured%20correctly,%20select

upvoted 3 times

HOTSPOT -

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE: Each correct selection is worth one point.

Hot Area:

### Answer Area

Statements	Yes	No
Platform as a service (PaaS) database offerings in Azure provide built-in high availability.	<input type="radio"/>	<input type="radio"/>
Platform as a service (PaaS) database offerings in Azure provide configurable scaling options.	<input type="radio"/>	<input type="radio"/>
Platform as a service (PaaS) database offerings in Azure reduce the administrative overhead for managing hardware.	<input type="radio"/>	<input type="radio"/>

### Answer Area

Statements	Yes	No
Platform as a service (PaaS) database offerings in Azure provide built-in high availability.	<input checked="" type="radio"/>	<input type="radio"/>
Platform as a service (PaaS) database offerings in Azure provide configurable scaling options.	<input checked="" type="radio"/>	<input type="radio"/>
Platform as a service (PaaS) database offerings in Azure reduce the administrative overhead for managing hardware.	<input checked="" type="radio"/>	<input type="radio"/>

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/sql-database-paas-overview>

 **tms2020** Highly Voted 3 years, 9 months ago

Yes Answer is correct

upvoted 61 times

 **Deliboi** Highly Voted 2 years, 2 months ago

Perfect Marketing LOL

upvoted 30 times

 **AGTraining** Most Recent 6 months, 2 weeks ago

Yes, Yes. Yes

upvoted 1 times

 **mamahani** 1 year, 5 months ago

I took the exam 12.04.2023 and got 1000, 100% score; this question appeared, correct answer:

yes / yes / yes

ps. perfect marketing lol

upvoted 11 times

 **david\_dlrio** 11 months, 1 week ago

More or less, which percentage of the questions here appeared in your exam?? Thanks :)

upvoted 1 times

 **chiiiweiii** 1 year, 6 months ago

Correct Answer

upvoted 3 times

 **oliverio** 2 years, 4 months ago

the perfect marketing :) YYY

upvoted 11 times

- 🗨️ **dk87** 2 years, 4 months ago  
Exam: 05/18/2022  
upvoted 5 times
- 🗨️ **D\_Mo\_ZY** 2 years, 6 months ago  
Answer is Correct  
Was in the exam (18/03/2022)  
upvoted 7 times
- 🗨️ **durak** 2 years, 10 months ago  
YES  
YES  
YES  
upvoted 8 times
- 🗨️ **jorgenoguerah** 3 years, 4 months ago  
The last one has a tricky word, It's said "Reduce" instead "Has NOT", PaaS eliminate the administrative overhead for managing hardware like IaaS and SaaS  
upvoted 10 times
- 🗨️ **okayhey** 6 months, 1 week ago  
It is reduce, because PaaS (and even SaaS) doesn't completely eliminate overhead costs for hardware, you still need computers to connect to Azure's VMs, printers, etc. You might also have a hybrid cloud, so you have to maintain your own servers  
upvoted 1 times
- 🗨️ **Average\_Joe** 3 years, 2 months ago  
Keyword there is 'managing hardware' bro  
upvoted 9 times
- 🗨️ **sdokmak** 2 years, 6 months ago  
has not the administrative overhead? doesn't make sense.  
upvoted 4 times
- 🗨️ **Tata11** 1 year, 11 months ago  
Dear Jorge: "Overhead" means operating costs. To say "has not administrative overhead" Doesn't make any sense.  
upvoted 5 times
- 🗨️ **ColoradoChick** 1 year, 2 months ago  
It seems obvious that jorgenoguerah meant "does not have". Not everyone speaks English as their 1st language.  
... and he's right. "Reduce" feels tricky, but hardware questions on the AZ-900 always mean physical hardware, and by "reduce", Microsoft means "reduces managing hardware vs on-prem solutions". MS doesn't always word it well in questions.  
upvoted 1 times
- 🗨️ **yeahno** 3 years, 5 months ago  
"Cloud features such as scalability, high-availability and multi-tenant capability are included, reducing the amount of coding that developers must do." from  
<https://azure.microsoft.com/en-gb/overview/what-is-paas/>  
upvoted 3 times
- 🗨️ **Billybob0604** 3 years, 5 months ago  
Only single databases I just read. "Single databases in Azure SQL Database support manual dynamic scalability, but not autoscale"  
upvoted 2 times
- 🗨️ **ANath** 3 years, 5 months ago  
1st one no doubt 'YES'.  
2nd one I am not sure because it says 'PaaS stands for Platform-as-a-service. Rather than creating a virtual infrastructure, and installing and managing the database software yourself, a PaaS solution does this for you. You specify the resources that you require (based on how large you think your databases will be, the number of users, and the performance you require), and Azure automatically creates the necessary virtual machines, networks, and other devices for you. You can usually scale up or down (increase or decrease the size and number of resources) quickly, as the volume of data and the amount of work being done varies; Azure handles this scaling for you, and you don't have to manually add or remove virtual machines, or perform any other form of configuration.'  
  
3rd one may be 'YES' but PaaS reduces administrative overhead of managing software. IaaS reduces overhead of managing Hardware.  
upvoted 6 times

🗨️ 👤 **Crazy\_Angel** 3 years, 2 months ago

I think 2nd one is yes too.

" Single databases in Azure SQL Database support manual dynamic scalability, but not autoscale. "

<https://docs.microsoft.com/en-us/azure/azure-sql/database/scale-resources>

upvoted 2 times

🗨️ 👤 **SitharthanMani** 3 years, 6 months ago

Second one, i need to check in portal

upvoted 2 times

HOTSPOT -

You have the following SQL query.

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

What are dbo.Products and ProductName? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

### Answer Area

dbo.Products:  ▼

A column
A database
A table
An index

ProductName:  ▼

A column
A database
A table
An index

### Answer Area

Suggested Answer:

dbo.Products:  ▼

A column
A database
A table
An index

ProductName:  ▼

A column
A database
A table
An index

 **AzureDP900** Highly Voted 2 years, 11 months ago

Table, column

upvoted 51 times

 **D\_Mo\_ZY** Highly Voted 2 years, 6 months ago

Correct

Was in the exam (18/03/2022)

upvoted 14 times

 **AGTraining** Most Recent 6 months, 2 weeks ago

Table,  
column

upvoted 1 times

 **98761234** 1 year, 5 months ago

Was in exam 02/04/2023

upvoted 4 times

 **chiiweiii** 1 year, 6 months ago

Correct Answer

upvoted 4 times

 **nmbuff** 1 year, 11 months ago

if you got fail this question, go back to study.

upvoted 7 times

🗨️ 👤 **Alex4190** 1 year, 6 months ago

Useless comment.

upvoted 10 times

🗨️ 👤 **Motivator** 2 years ago

Table

Column

upvoted 5 times

🗨️ 👤 **Eltooth** 2 years, 2 months ago

Table

Column

upvoted 4 times

🗨️ 👤 **w1920** 2 years, 6 months ago

table, column

upvoted 6 times

🗨️ 👤 **Yvekke** 2 years, 8 months ago

Table, column

upvoted 5 times

🗨️ 👤 **Shams132** 2 years, 9 months ago

Great Question lol

upvoted 2 times

🗨️ 👤 **Learner\_Thunder2** 2 years, 9 months ago

How can an official certification exam hold this kind of questions?

upvoted 4 times

🗨️ 👤 **vp45** 2 years, 9 months ago

I am not complaining, if i get few of these.

upvoted 10 times

🗨️ 👤 **Aneta\_ane** 2 years, 9 months ago

Came up on my exam today.

upvoted 4 times

🗨️ 👤 **azurebilearning** 2 years, 11 months ago

correct

upvoted 1 times

🗨️ 👤 **bettz\_druid** 3 years ago

correct!

upvoted 3 times

🗨️ 👤 **dumbmf** 3 years, 1 month ago

It is not wrong

upvoted 2 times

🗨️ 👤 **ANath** 3 years, 5 months ago

correct

upvoted 4 times

HOTSPOT -

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE: Each correct selection is worth one point.

Hot Area:

## Answer Area

Statements	Yes	No
You must apply operating system updates to Azure SQL databases regularly.	<input type="radio"/>	<input type="radio"/>
You need a Microsoft 365 subscription to create an Azure SQL database.	<input type="radio"/>	<input type="radio"/>
You can use existing Microsoft SQL Server licenses to reduce the cost of Azure SQL databases.	<input type="radio"/>	<input type="radio"/>

## Answer Area

Statements	Yes	No
<b>Suggested Answer:</b> You must apply operating system updates to Azure SQL databases regularly.	<input type="radio"/>	<input checked="" type="radio"/>
You need a Microsoft 365 subscription to create an Azure SQL database.	<input type="radio"/>	<input checked="" type="radio"/>
You can use existing Microsoft SQL Server licenses to reduce the cost of Azure SQL databases.	<input checked="" type="radio"/>	<input type="radio"/>

Reference:

<https://azure.microsoft.com/en-gb/blog/hot-patching-sql-server-engine-in-azure-sql-database/> <https://azure.microsoft.com/en-us/services/sql-database/#product-overview>

 **AzureDP900** Highly Voted 2 years, 11 months ago

No, No, Yes

upvoted 57 times

 **gme999** Highly Voted 2 years, 11 months ago

3rd one is Yes. <https://azure.microsoft.com/en-us/pricing/hybrid-benefit/faq/>

Azure Hybrid Benefit allows you to use SQL Server licenses with Software Assurance or qualifying subscription licenses to pay a reduced base rate\* for these products and services for SQL Server on Azure:

vCPU-based service tiers of Azure SQL Database (excluding serverless).

Azure SQL Managed Instance.

SQL Server in Azure Virtual Machines.

SQL Server Integration Services.

upvoted 20 times

 **rajat7878** 2 years, 6 months ago

Agreed. Third one is Yes.

upvoted 4 times

 **sjokkorein** Most Recent 6 months ago

Classic question:

You must = Almost certain "NO"

You need = Almost certain "NO"

upvoted 2 times

🗨️ 👤 **AGTraining** 6 months, 2 weeks ago

No, No, Yes

upvoted 1 times

🗨️ 👤 **Oluniyi** 8 months, 3 weeks ago

The last one should be N. For PaaS, it is N.

<https://learn.microsoft.com/en-us/azure/azure-sql/azure-hybrid-benefit?view=azuresql&tabs=azure-powershell>

How does Azure Hybrid Benefit for SQL Server differ from license mobility?

We offer license mobility benefits to SQL Server customers with Software Assurance. License mobility allows reassignment of their licenses to a partner's shared servers. You can use this benefit on Azure IaaS and AWS EC2.

Azure Hybrid Benefit for SQL Server differs from license mobility in two key areas:

It provides economic benefits for moving highly virtualized workloads to Azure. SQL Server Enterprise Edition customers can get four cores in Azure in the General Purpose SKU for every core they own on-premises for highly virtualized applications. License mobility doesn't allow any special cost benefits for moving virtualized workloads to the cloud.

upvoted 2 times

🗨️ 👤 **chiiweiii** 1 year, 6 months ago

NNY are Correct Answer

upvoted 4 times

🗨️ 👤 **raydel92** 1 year, 11 months ago

No, No, Yes

<https://azure.microsoft.com/en-us/pricing/hybrid-benefit/#why-azure-hybrid-benefit>

upvoted 4 times

🗨️ 👤 **WowThato** 2 years, 1 month ago

Feels like the perfect way to market yourself

upvoted 6 times

🗨️ 👤 **Eltooth** 2 years, 2 months ago

NO

NO

Yes

upvoted 3 times

🗨️ 👤 **heenak55** 2 years, 4 months ago

In exam 21 April 2022

upvoted 4 times

🗨️ 👤 **RichMelbourne** 2 years, 5 months ago

Shown on test 15/04/2022

upvoted 3 times

🗨️ 👤 **SiDoCiOuS** 2 years, 5 months ago

Appeared on the exam on 03/22/2022

Answers are correct! No, No, Yes.

upvoted 6 times

🗨️ 👤 **wando5000** 2 years, 9 months ago

Currysing you are incorrect as the first 2 lines do testify (Azure SQL and SQL managed instance) - so its NOT just SQL installed on a vm

vCPU-based service tiers of Azure SQL Database (excluding serverless).

Azure SQL Managed Instance

upvoted 2 times

🗨️ 👤 **Sam\_Zee** 2 years, 9 months ago

N/N/Y is the correct answer

upvoted 3 times

🗨️ 👤 **theLazyGuy** 2 years, 10 months ago

N/N/Y

Can I bring my on-premises or IaaS SQL Server license to Hyperscale? Yes, Azure Hybrid Benefit ... (<https://docs.microsoft.com/en-us/azure/azure-sql/database/service-tier-hyperscale-frequently-asked-questions-faq>  
<https://azure.microsoft.com/en-us/pricing/hybrid-benefit/#calculator>)

upvoted 4 times

🗨️ 👤 **georgich87** 2 years, 11 months ago

<https://azure.microsoft.com/en-in/blog/eight-ways-to-optimize-costs-on-azure-sql/>

upvoted 2 times

🗨️ 👤 **gme999** 2 years, 11 months ago

Azure Hybrid Benefit allows you to use SQL Server licenses with Software Assurance or qualifying subscription licenses to pay a reduced base rate\* for these products and services for SQL Server on Azure:

vCPU-based service tiers of Azure SQL Database (excluding serverless).

Azure SQL Managed Instance.

SQL Server in Azure Virtual Machines.

SQL Server Integration Services.

upvoted 5 times

Which statement is an example of Data Definition Language (DDL)?

- A. SELECT
- B. JOIN
- C. UPDATE
- D. CREATE

**Suggested Answer:** D

Data Definition Language (DDL) statements defines data structures. Use these statements to create, alter, or drop data structures in a database. These statements include:

- ⇒ ALTER
- ⇒ Collations
- ⇒ CREATE
- ⇒ DROP
- ⇒ DISABLE TRIGGER
- ⇒ ENABLE TRIGGER
- ⇒ RENAME
- ⇒ UPDATE STATISTICS
- ⇒ TRUNCATE TABLE
- ⇒ INSERT

Reference:

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

Community vote distribution

D (100%)

🗳️ **AzureDP900** Highly Voted 2 years, 11 months ago

Create is correct D  
upvoted 30 times

🗳️ **Pora2000** Highly Voted 2 years, 4 months ago

Answer lists (explanation) INSERT as one. That is DML  
upvoted 13 times

🗳️ **AGTraining** Most Recent 6 months, 2 weeks ago

create  
upvoted 1 times

🗳️ **akhil5432** 1 year, 1 month ago

Selected Answer: D  
create  
upvoted 2 times

🗳️ **vicks60** 1 year, 5 months ago

Selected Answer: D  
CREATE is correct  
upvoted 3 times

🗳️ **chiiiweiii** 1 year, 6 months ago

Selected Answer: D  
Correct Answer  
upvoted 3 times

🗳️ **iesRuben** 1 year, 9 months ago

Selected Answer: D  
CREATE is the answer  
upvoted 3 times

🗨️ **JA2018** 1 year, 11 months ago

**Selected Answer: D**

part of DDL

upvoted 3 times

🗨️ **PatYeo** 1 year, 11 months ago

D is the correct Answer.

1. DDL

- CREATE

- ALTER

- DROP

2. DCL

- GRANT

- REVOKE

3. DML

- INSERT

- DELETE

- UPDATE

- SELECT

upvoted 8 times

🗨️ **Nurialzard** 1 year, 7 months ago

Select is DQL

upvoted 1 times

🗨️ **Aziza\_Adam** 1 year, 4 months ago

Plus merge

upvoted 1 times

🗨️ **WSW\_IT** 2 years ago

Not only INSERT is DML, but also TRUNCATE TABLE since it doesn't change the structure, only data.

upvoted 1 times

🗨️ **icew3b** 2 years, 1 month ago

FYI: in the Text is says INSERT is part of DEFINITION. that's wrong.... INSERT is manipulation

upvoted 2 times

🗨️ **cp2323** 2 years, 4 months ago

Answer is D but why did you include INSERT in DDL on your answer area? Please remove that

upvoted 5 times

🗨️ **SiDoCiOuS** 2 years, 5 months ago

**Selected Answer: D**

Appeared on the exam on 03/22/2022

Answer is correct!

upvoted 3 times

🗨️ **Suyu** 2 years, 6 months ago

CREATE for sure

upvoted 2 times

🗨️ **G\_unit\_19** 2 years, 7 months ago

Still CREATE is the correct answer

upvoted 1 times

🗨️ **Loofy** 2 years, 11 months ago

For sure is CREATE! -D

upvoted 2 times

HOTSPOT -

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE: Each correct selection is worth one point.

Hot Area:

### Answer Area

Statements	Yes	No
Azure Data Studio can be used to query an Azure SQL database from a device that runs macOS.	<input type="radio"/>	<input type="radio"/>
Microsoft SQL Server Management Studio (SSMS) enables users to create and use SQL notebooks.	<input type="radio"/>	<input type="radio"/>
Azure Data Studio can be used to restore a database.	<input type="radio"/>	<input type="radio"/>

Suggested Answer:

### Answer Area

Statements	Yes	No
Azure Data Studio can be used to query an Azure SQL database from a device that runs macOS.	<input checked="" type="radio"/>	<input type="radio"/>
Microsoft SQL Server Management Studio (SSMS) enables users to create and use SQL notebooks.	<input type="radio"/>	<input checked="" type="radio"/>
Azure Data Studio can be used to restore a database.	<input checked="" type="radio"/>	<input type="radio"/>

Box 1: Yes -

Azure Data Studio is a cross-platform database tool for data professionals using on-premises and cloud data platforms on Windows, macOS, and Linux.

You can use Azure Data Studio to connect to an Azure SQL Database server. You'll then run Transact-SQL (T-SQL) statements to create and query Azure SQL databases.

Box 2: No -

SQL Server Management Studio is for configuring, managing, and administering all components within Microsoft SQL Server, not to create SQL notebooks.

Instead use Azure Data Studio to create SQL notebook.

Box 3: Yes -

You can use the Azure Data Studio to restore databases.

Reference:

<https://docs.microsoft.com/en-us/sql/azure-data-studio/what-is-azure-data-studio>

 **dayan911** Highly Voted 3 years, 2 months ago

1. Yes
- 2.No
- 3.Yes

upvoted 58 times

 **Meyti** Highly Voted 2 years, 9 months ago

I was not sure with option 3 , but the answer is correct, and the source:

<https://docs.microsoft.com/en-us/sql/azure-data-studio/what-is-azure-data-studio?view=sql-server-ver15#database-administration>

upvoted 16 times

🗨️ 👤 **Jangoman** 2 years, 7 months ago

You can restore but cant take backups

upvoted 6 times

🗨️ 👤 **BabyBee** 2 years, 5 months ago

Can do Restore/Backups with Azure Data Studio.

Feature Azure Data Studio SSMS

Backup / Restore Yes Yes

upvoted 4 times

🗨️ 👤 **PI\_Team** Most Recent 4 months ago

see the tables!

BTW, option 3 is in preview so I presume Yes is the correct answer

<https://learn.microsoft.com/en-us/azure-data-studio/what-is-azure-data-studio?view=sql-server-ver16>

upvoted 2 times

🗨️ 👤 **AGTraining** 6 months, 2 weeks ago

1. Yes

2.No

3.Yes

upvoted 1 times

🗨️ 👤 **gabo** 7 months ago

Backup/Restore option is in Preview mode in Azure Data Studio, but never the less, it's available, so the answer is YES.

upvoted 2 times

🗨️ 👤 **jxs221** 7 months ago

The correct answer is Yes , No , NO

upvoted 1 times

🗨️ 👤 **chiiiweiii** 1 year, 6 months ago

YNY are Correct Answer

upvoted 3 times

🗨️ 👤 **xRiot007** 2 years, 2 months ago

Yes - you can use Az Data Studio to query DBs. The device is irrelevant

No

Yes - the tool can do restores.

upvoted 5 times

🗨️ 👤 **RichMelbourne** 2 years, 5 months ago

Shown on test 15/04/2022

upvoted 5 times

🗨️ 👤 **SiDoCiOuS** 2 years, 5 months ago

Appeared on the exam on 03/22/2022

Answers are correct! Yes, No, Yes

upvoted 5 times

🗨️ 👤 **Meyti** 2 years, 9 months ago

I was not sure with option 3 , but the answer is correct, and the source:

<https://docs.microsoft.com/en-us/sql/azure-data-studio/what-is-azure-data-studio?view=sql-server-ver15#database-administration>

upvoted 2 times

🗨️ 👤 **azurebilearning** 2 years, 11 months ago

Correct Answers

upvoted 2 times

🗨️ 👤 **MONALISA10** 2 years, 12 months ago

answers are correct

upvoted 1 times

🗨️ 👤 **valente\_sven1** 3 years ago

Given answers are correct  
upvoted 1 times

🗨️ 👤 **azureprac1** 3 years, 2 months ago

Option 1 seem to be incorrect  
upvoted 2 times

🗨️ 👤 **NarendraGillalu** 3 years, 2 months ago

Correct!  
upvoted 3 times

🗨️ 👤 **lhueghian** 3 years, 3 months ago

SQL notebook: Azure Data Studio instead  
upvoted 6 times

You are deploying a software as a service (SaaS) application that requires a relational database for Online Transaction Processing (OLTP). Which Azure service should you use to support the application?

- A. Azure Cosmos DB
- B. Azure HDInsight
- C. Azure SQL Database
- D. Azure Synapse Analytics

**Suggested Answer: C**

Azure SQL Database is relational database and a managed service.

Incorrect Answers:

A, B: Cosmos DB, HDInsight are non-relational databases.

D: Azure Synapse Analytics is for data warehousing, not for Online Transaction Processing

Reference:

<https://cloud.netapp.com/blog/azure-cvo-blg-azure-database-review-your-guide-for-database-assessment>

Community vote distribution

C (100%)

-  **yeanlingmedal71** Highly Voted 3 years, 3 months ago  
Correct answer and explanation  
upvoted 33 times
-  **azurebilearning** Highly Voted 2 years, 11 months ago  
Azure SQL Database is correct answer for OLTP  
upvoted 15 times
-  **AGTraining** Most Recent 6 months, 2 weeks ago  
option C  
upvoted 1 times
-  **m0j7zdg** 7 months, 2 weeks ago  
Azure SQL DB is PaaS  
upvoted 1 times
-  **AZFabio** 1 year ago  
Selected Answer: C  
SQL Database  
upvoted 1 times
-  **DaliborUhlir** 1 year, 1 month ago  
Selected Answer: C  
Correct is C, Azure SQL DB  
upvoted 1 times
-  **mamahani** 1 year, 5 months ago  
I took the exam 12.04.2023; this question appeared, correct answer:  
Azure SQL database  
upvoted 6 times
-  **paulk97** 1 year, 6 months ago  
Azure SLQ DB is PaaS, not SaaS.  
upvoted 5 times
-  **chiiiweiii** 1 year, 6 months ago  
Selected Answer: C  
Correct Answer  
upvoted 1 times

- 🗨️ 👤 **Edriv** 1 year, 9 months ago  
If it is OLAP would it be AZURE COSMOS DB? by the way I choose C  
upvoted 1 times
- 🗨️ 👤 **nmbuff** 1 year, 11 months ago  
**Selected Answer: C**  
C is correct answer.  
upvoted 3 times
- 🗨️ 👤 **Breeze21** 2 years, 2 months ago  
Exam: 19.07.22  
upvoted 2 times
- 🗨️ 👤 **Eltooth** 2 years, 2 months ago  
**Selected Answer: C**  
C is correct answer.  
upvoted 2 times
- 🗨️ 👤 **dk87** 2 years, 4 months ago  
Exam: 05/18/2022  
upvoted 3 times
- 🗨️ 👤 **bfan95** 2 years, 4 months ago  
**Selected Answer: C**  
Answer: C  
upvoted 2 times
- 🗨️ 👤 **Vishnu007** 2 years, 5 months ago  
YES IT IS  
upvoted 1 times
- 🗨️ 👤 **D\_Mo\_ZY** 2 years, 6 months ago  
Correct  
Was in the exam (18/03/2022)  
upvoted 5 times

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

**Suggested Answer:** AD

A: 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.

D: 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.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/sql-database-paas-overview>

Community vote distribution

AD (100%)

-  **AleTato** Highly Voted 3 years, 5 months ago  
AD - Correct  
upvoted 64 times
-  **PhyoPluto** Highly Voted 1 year, 6 months ago  
the exam on 07/03/2023  
upvoted 10 times
-  **AGTraining** Most Recent 6 months, 2 weeks ago  
option A e D  
upvoted 1 times
-  **VNR25** 1 year, 2 months ago  
This question appeared in the 07/07/23 exam and the answer is correct.  
upvoted 2 times
-  **chiiiweiii** 1 year, 6 months ago  
Selected Answer: AD  
Correct Answers  
upvoted 2 times
-  **JA2018** 1 year, 11 months ago  
Selected Answer: AD  
PaaS manages the regular backups and software updates.  
upvoted 3 times
-  **Dasan** 2 years, 1 month ago  
Correct  
upvoted 2 times
-  **johnsmith11235** 2 years, 1 month ago  
A & D - Correct  
upvoted 3 times
-  **Eltooth** 2 years, 2 months ago  
Selected Answer: AD  
A and D are correct answers.

upvoted 1 times

🗨️ 👤 **Gautam123** 2 years, 3 months ago

Selected Answer: AD

Correct

upvoted 1 times

🗨️ 👤 **heenak55** 2 years, 4 months ago

In exam 21 April 2022

upvoted 4 times

🗨️ 👤 **RichMelbourne** 2 years, 5 months ago

Shown on test 15/04/2022

upvoted 4 times

🗨️ 👤 **SiDoCiOuS** 2 years, 5 months ago

Selected Answer: AD

Appeared on the exam on 03/22/2022

Answers are correct!

upvoted 4 times

🗨️ 👤 **xiban** 2 years, 11 months ago

why not B ?

upvoted 2 times

🗨️ 👤 **crainos** 2 years, 10 months ago

because paas backup and update automaticaly by microsoft

upvoted 9 times

🗨️ 👤 **sethuramansp** 2 years, 6 months ago

Let's say if the option is tweaked a bit like "fully managed backup n restore processes" then it matches PAAS, but the statement is complete control which is sort of an User or Administrator having control over those processes which is against the existence of PAAS features on a whole.

upvoted 9 times

🗨️ 👤 **JA2018** 1 year, 11 months ago

Complete control will require IaaS

upvoted 1 times

🗨️ 👤 **AzureDP900** 2 years, 11 months ago

A & D is correct answer

upvoted 3 times

🗨️ 👤 **smario** 2 years, 11 months ago

Yes correct

upvoted 2 times

🗨️ 👤 **MONALISA10** 2 years, 12 months ago

correct

upvoted 2 times

HOTSPOT -

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE: Each correct selection is worth one point.

Hot Area:

### 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.	<input type="radio"/>	<input type="radio"/>
If you have a platform as a service (PaaS) database in Azure, backups are performed automatically.	<input type="radio"/>	<input type="radio"/>
If you have a platform as a service (PaaS) database in Azure, you are responsible for installation of the database engine.	<input type="radio"/>	<input type="radio"/>

Suggested Answer:

### 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.	<input type="radio"/>	<input checked="" type="radio"/>
If you have a platform as a service (PaaS) database in Azure, backups are performed automatically.	<input checked="" type="radio"/>	<input type="radio"/>
If you have a platform as a service (PaaS) database in Azure, you are responsible for installation of the database engine.	<input type="radio"/>	<input checked="" type="radio"/>

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 -

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/sql-database-paas-overview>

 **svetlanam** Highly Voted 2 years, 4 months ago

Correct. No, Yes, No  
upvoted 36 times

 **studysmart** Highly Voted 1 year, 11 months ago

Confirmed, the answer is No, Yes, No  
It was in 26/09/2022 exam and I scored 1000.  
upvoted 23 times

 **AGTraining** Most Recent 6 months, 2 weeks ago

No,  
Yes,  
No  
upvoted 1 times

 **chiiiweiii** 1 year, 6 months ago

NYN are Correct Answer

upvoted 3 times

  **NafiDre** 1 year, 10 months ago

No

Yes

No

upvoted 4 times

  **Breeze21** 2 years, 2 months ago

Exam: 19.07.22

upvoted 4 times

  **Eltooth** 2 years, 2 months ago

No

Yes

No

upvoted 3 times

  **RakeshSripalle** 2 years, 2 months ago

No, Yes, No is correct.

upvoted 3 times

  **simplementeluca** 2 years, 2 months ago

No, Yes, No ..

upvoted 3 times

  **bfan95** 2 years, 4 months ago

N,Y,N is correct

upvoted 4 times

  **OPT\_001122** 2 years, 4 months ago

correct answer

upvoted 4 times

DRAG DROP -

You have a table named Sales that contains the following data.

SalesDate	SalesAmount	ProductID
4-Apr-20	\$2,000	1
5-Apr-20	\$40	2
5-Apr-20	\$2,300	1
6-Apr-20	\$40	3
6-Apr-20	\$200	4

You need to query the table to return the average sales amount per day. The output must produce the following results.

SalesDate	AVG(SalesAmount)
4-Apr-20	\$2,000
5-Apr-20	\$1,170
6-Apr-20	\$120

How should you complete the query? To answer, drag the appropriate values to the correct targets. Each value 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:

#### Values

CREATE	GROUP BY
ORDER BY	SELECT

#### Answer Area

Value	SalesDate, AVG(SalesAmount)
FROM Sales	
Value	SalesDate
ORDER BY SalesDate	

#### Suggested Answer:

##### Values

CREATE	GROUP BY
ORDER BY	SELECT

##### Answer Area

SELECT	SalesDate, AVG(SalesAmount)
FROM Sales	
GROUP BY	SalesDate
ORDER BY SalesDate	

Box 1: SELECT -

Box 2: GROUP BY -

Example:

When used with a GROUP BY clause, each aggregate function produces a single value covering each group, instead of a single value covering the whole table.

The following example produces summary values for each sales territory in the AdventureWorks2012 database. The summary lists the average bonus received by the sales people in each territory, and the sum of year-to-date sales for each territory.

```
SELECT TerritoryID, AVG(Bonus)as 'Average bonus', SUM(SalesYTD) as 'YTD sales'
```

```
FROM Sales.SalesPerson -
```

```
GROUP BY TerritoryID;
```

Reference:

<https://docs.microsoft.com/en-us/sql/t-sql/functions/avg-transact-sql>

- 🗨️ 👤 **AleTato** Highly Voted 3 years, 5 months ago  
LOL! It's real this question?  
Anyway the answer is CORRECT.. but.... LOOOOOOLLLLL  
upvoted 71 times
- 🗨️ 👤 **AnonymousJhb** 1 year, 11 months ago  
Don't be a buttstick. This is not a sql exam, it's a data engineering exam. There are prob non sql ppl here.  
upvoted 31 times
- 🗨️ 👤 **Tigereyes19** 2 years, 11 months ago  
hahaha  
upvoted 3 times
- 🗨️ 👤 **mattzocsgogod** 2 years, 10 months ago  
yesnt kekw  
upvoted 4 times
- 🗨️ 👤 **SilNilanjan** 3 years, 3 months ago  
This exam has got non technical audiences as well :)  
upvoted 36 times
- 🗨️ 👤 **sdokmak** 2 years, 6 months ago  
or people that never touched sql before and want to smash this exam solely with the help of exam topics (me lol).  
upvoted 29 times
- 🗨️ 👤 **sandyshd** Highly Voted 3 years, 3 months ago  
correct  
upvoted 10 times
- 🗨️ 👤 **AGTraining** Most Recent 6 months, 2 weeks ago  
SELECT  
GROUP BY  
upvoted 2 times
- 🗨️ 👤 **mamahani** 1 year, 5 months ago  
I took the exam 12.04.2023; this question appeared, correct answer:  
select / group by  
upvoted 6 times
- 🗨️ 👤 **chiiiweiii** 1 year, 6 months ago  
Correct Answer  
upvoted 1 times
- 🗨️ 👤 **AndreLima** 1 year, 11 months ago  
Avanced level to this exam DP-900....  
But...but lets go !  
upvoted 2 times
- 🗨️ 👤 **xRiot007** 2 years, 2 months ago  
SELECT - to have any results, you need a SELECT statement.  
GROUP BY - The sales date is our criteria of aggregating the rows.  
Because there is an AVG() function call in the SELECT, the resulting groups will have 1 row each, resulting in the table below  
upvoted 9 times
- 🗨️ 👤 **Eltooth** 2 years, 2 months ago  
Select  
Group By  
upvoted 2 times
- 🗨️ 👤 **dk87** 2 years, 4 months ago  
Exam: 05/18/2022  
upvoted 4 times

🗨️ 👤 **RichMelbourne** 2 years, 5 months ago

Shown on test 15/04/2022

upvoted 4 times

🗨️ 👤 **SiDoCiOuS** 2 years, 5 months ago

Appeared on the exam on 03/22/2022

Answers are correct! Select/Group by

upvoted 5 times

🗨️ 👤 **Yvekke** 2 years, 8 months ago

correct

upvoted 1 times

🗨️ 👤 **cianryan** 2 years, 10 months ago

[SELECT], [GROPU BY] is the nswer.

upvoted 2 times

🗨️ 👤 **ahee** 2 years, 10 months ago

SELECT & GROUP BY is the correct answer.

upvoted 2 times

🗨️ 👤 **Drummer** 3 years, 5 months ago

The answer is correct!

upvoted 2 times

When you create an Azure SQL database, which account can always connect to the database?

- A. the Azure Active Directory (Azure AD) account that created the database
- B. the server admin login account of the logical server
- C. the Azure Active Directory (Azure AD) administrator account
- D. the sa account

**Suggested Answer: B**

When you first deploy Azure SQL, you specify an admin login and an associated password for that login. This administrative account is called Server admin.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/single-database-create-quickstart>

Community vote distribution

B (100%)

 **faiqsr1** Highly Voted 3 years, 4 months ago

Answer is B. When you first deploy Azure SQL, you specify an admin login and an associated password for that login. This administrative account is called Server admin. Link :<https://docs.microsoft.com/en-us/azure/azure-sql/database/logins-create-manage>  
upvoted 70 times

 **peterp999** Highly Voted 3 years, 5 months ago

B. Verified in AZ portal.  
upvoted 26 times

 **AGTraining** Most Recent 6 months, 2 weeks ago

option B  
upvoted 1 times

 **gabo** 7 months ago

Answer is correct, provided the admin user is trying to connect from an IP that is allowed through the DB firewall.  
upvoted 1 times

 **akhil5432** 1 year, 1 month ago

Selected Answer: B  
OPTION B  
upvoted 1 times

 **chiiiweiii** 1 year, 6 months ago

Selected Answer: B  
Correct Answer  
upvoted 2 times

 **Stuurwiel1** 1 year, 7 months ago

Selected Answer: B  
When you create an Azure SQL database, the server admin login account of the logical server, also known as the "server admin" or "administrator" account, can always connect to the database. This account is created when the logical server is created and is used to manage the server and its databases.  
upvoted 3 times

 **Pinha** 1 year, 7 months ago

B. the server admin login account of the logical server can always connect to the database. This account is also known as the "server-level principal login" and it is created automatically when you create a new Azure SQL server. This account has the highest level of permissions and can be used to perform all management tasks on the server and the databases it contains.  
upvoted 1 times

 **Edriv** 1 year, 10 months ago

Hmmm... Maybe D  
upvoted 1 times

🗨️ 👤 **Eltooth** 2 years, 2 months ago

**Selected Answer: B**

B is correct answer.

upvoted 3 times

🗨️ 👤 **podvad** 2 years, 4 months ago

why not D?

upvoted 3 times

🗨️ 👤 **bfan95** 2 years, 4 months ago

**Selected Answer: B**

Answer: B

upvoted 3 times

🗨️ 👤 **RichMelbourne** 2 years, 5 months ago

Shown on test 15/04/2022

upvoted 3 times

🗨️ 👤 **SiDoCiOuS** 2 years, 5 months ago

**Selected Answer: B**

Appeared on the exam on 03/22/2022

Answer is correct!

upvoted 2 times

🗨️ 👤 **D\_Mo\_ZY** 2 years, 6 months ago

B.

Was in the exam (18/03/2022)

upvoted 4 times

🗨️ 👤 **Koph** 2 years, 9 months ago

Server admin is correct

upvoted 1 times

🗨️ 👤 **topeia** 2 years, 11 months ago

B is correct

upvoted 1 times

Which statement is an example of Data Definition Language (DDL)?

- A. SELECT
- B. INSERT
- C. DELETE
- D. DROP

**Suggested Answer:** D

Data Definition Language (DDL) statements defines data structures. Use these statements to create, alter, or drop data structures in a database. These statements include:

- ⇒ ALTER
- ⇒ Collations
- ⇒ CREATE
- ⇒ DROP
- ⇒ DISABLE TRIGGER
- ⇒ ENABLE TRIGGER
- ⇒ RENAME
- ⇒ UPDATE STATISTICS
- ⇒ TRUNCATE TABLE
- ⇒ UPDATE

Reference:

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

Community vote distribution

D (100%)

 **Aeditya** Highly Voted 2 years, 9 months ago

correct answer option D  
upvoted 32 times

 **cp2323** Highly Voted 2 years, 4 months ago

Remove UPDATE from your answer explanation area as its not DDL  
upvoted 15 times

 **AGTraining** Most Recent 6 months, 2 weeks ago

option D  
upvoted 1 times

 **DataEngDP** 1 year, 4 months ago

UPDATE is NOT a DDL statement, please delete it from the explanation as that is DML statement.  
upvoted 1 times

 **Vanessa23** 1 year, 5 months ago

DROP is definition keyword  
upvoted 2 times

 **chiiiweiii** 1 year, 6 months ago

Selected Answer: D  
Correct Answer  
upvoted 1 times

 **sreekan2** 1 year, 8 months ago

Selected Answer: D  
correct answer option D  
upvoted 1 times

 **sreekan2** 1 year, 8 months ago

Selected Answer: D

correct answer option D

upvoted 1 times

🗨️ 👤 **JA2018** 1 year, 11 months ago

**Selected Answer: D**

Part of DDL

upvoted 2 times

🗨️ 👤 **JA2018** 1 year, 11 months ago

**Selected Answer: D**

part of DDL which is required for this operation

upvoted 2 times

🗨️ 👤 **Breeze21** 2 years, 2 months ago

Exam: 19.07.22

upvoted 2 times

🗨️ 👤 **LaloSalamanca** 2 years, 2 months ago

it should be D

upvoted 1 times

🗨️ 👤 **RakeshSripalle** 2 years, 2 months ago

Remove UPDATE from answer explanation area as it is DML not DDL

upvoted 3 times

🗨️ 👤 **Manoj6** 2 years, 3 months ago

DROP is the answer

upvoted 3 times

🗨️ 👤 **Loofy** 2 years, 11 months ago

DROP is the - Answer

upvoted 5 times

🗨️ 👤 **Claudio\_Ale** 2 years, 11 months ago

The correct answer is "D". It's OK.

upvoted 3 times

A team of developers has computers that run Windows 10 and Ubuntu Desktop.

The developers need to connect to and query an Azure SQL database from each of their computers. The developers require code assistance features such as IntelliSense.

What should the developers use?

- A. sqlcmd
- B. Microsoft SQL Server Management Studio (SSMS)
- C. Azure Data Studio
- D. Azure Data Explorer

**Suggested Answer: C**

Azure Data Studio is a cross-platform database tool for data professionals who use on-premises and cloud data platforms on Windows, macOS, and Linux.

Azure Data Studio offers a modern editor experience with IntelliSense, code snippets, source control integration, and an integrated terminal.

Reference:

<https://docs.microsoft.com/en-us/sql/azure-data-studio/download-azure-data-studio>

Community vote distribution

C (100%)

 **AzureDP900** Highly Voted 2 years, 11 months ago

Correct Azure Data Studio  
upvoted 36 times

 **kukushka** Highly Voted 2 years, 9 months ago

It is not Microsoft SQL Server Management Studio (SSMS) because it would not be supported on Ubuntu.  
upvoted 12 times

 **jiggybrain** 2 years, 9 months ago

SSMS doesn't run queries, it handles configuration and management.  
upvoted 2 times

 **Seansmyrke** 2 years, 8 months ago

you can certainly run queries in SSMS the above comment is correct  
upvoted 14 times

 **gerryboy** 2 years, 6 months ago

ssms can run queries  
upvoted 2 times

 **AGTraining** Most Recent 6 months, 2 weeks ago

option C  
upvoted 1 times

 **Jejemon** 1 year, 3 months ago

Correct. Ubuntu is the modern, open source operating system on LINUX for the enterprise server, desktop, cloud, and IoT.  
upvoted 1 times

 **Prat25200607** 1 year, 6 months ago

Selected Answer: C

C is correct  
upvoted 3 times

 **chiiweiii** 1 year, 6 months ago

Selected Answer: C

Correct Answer  
upvoted 3 times

🗨️ 👤 **\_Jene\_** 1 year, 7 months ago

Why not Azure Data Explorer?

upvoted 2 times

🗨️ 👤 **Yomzie** 6 months, 3 weeks ago

Azure Data Explorer is a standalone service offering ON-DEMAND high-performance analysis of log & telemetry data (similar to how Azure Synapse Data Explorer functions).

upvoted 1 times

🗨️ 👤 **iesRuben** 1 year, 9 months ago

**Selected Answer: C**

Azure data studio is correct!

upvoted 3 times

🗨️ 👤 **raydel92** 1 year, 11 months ago

**Selected Answer: C**

Azure Data Studio

upvoted 3 times

🗨️ 👤 **JA2018** 1 year, 11 months ago

**Selected Answer: C**

Only viable option for Ubuntu users

upvoted 4 times

🗨️ 👤 **CrazyAunt** 2 years, 6 months ago

SSMS can run queries with intelisense, but it does not work on Ubuntu.

upvoted 9 times

🗨️ 👤 **gerryboy** 2 years, 6 months ago

ssms can run queries

upvoted 1 times

🗨️ 👤 **Koph** 2 years, 9 months ago

Azure Data Studio is correct

upvoted 3 times

HOTSPOT -

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

Hot Area:

### 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.

Suggested Answer:

### 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.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/transparent-data-encryption-tde-overview?tabs=azure-portal>

 **starkwinterfell** Highly Voted 3 years, 5 months ago  
Answer is correct

Transparent data encryption (TDE) helps protect Azure SQL Database, Azure SQL Managed Instance, and Azure Synapse Analytics against the threat of malicious offline activity by encrypting data at rest.  
upvoted 97 times

 **Eltooth** Highly Voted 2 years, 2 months ago  
Database. Always.  
upvoted 13 times

 **AGTraining** Most Recent 6 months, 2 weeks ago  
option C  
upvoted 1 times

 **jaszi** 12 months ago  
Transparent Data Encryption encrypts the database or the server?  
ChatGPT  
Transparent Data Encryption (TDE) encrypts the entire database, including the data files, log files, and backup files. It does not encrypt the server itself. TDE operates at the database level and provides encryption for data at rest. This means that the data is encrypted on the storage level, ensuring that even if someone gains unauthorized access to the physical storage media, they would not be able to read the data without the appropriate encryption key.  
upvoted 1 times

 **akhil5432** 1 year, 1 month ago  
IF THIS olatform is free why they are asking to PAY after 126 question.they are allowing to 126 quest. only  
upvoted 6 times

 **iVath** 1 year, 4 months ago  
C correct.  
pls ref to : <https://learn.microsoft.com/en-us/sql/relational-databases/security/encryption/transparent-data-encryption?view=sql-server-ver15>

where lines as :

One solution is to encrypt sensitive data in a database

About TDE

Encryption of a database file is done at the page level. The pages in an encrypted database .....

upvoted 1 times

🗨️ **chiiiweiii** 1 year, 6 months ago

Correct Answer

upvoted 3 times

🗨️ **Dusica** 1 year, 10 months ago

MS docs: "Transparent data encryption (TDE) encrypts SQL Server, Azure SQL Database, and Azure Synapse Analytics data files. This encryption is known as encrypting data at rest", The highlight is on data files; implemented on servers>inherited by databases which hold data that is encrypted. C is correct

upvoted 4 times

🗨️ **Gavin1272** 1 year, 11 months ago

I think what the MS Documentation is trying to say about Transparent Data Encryption (TDE) is that TDE is installed at the server level but only applied at the Database level. So, the server is not encrypted but the databases on the server are encrypted.

upvoted 3 times

🗨️ **simplementeluca** 2 years, 2 months ago

allways de database

upvoted 4 times

🗨️ **vkp17in** 2 years, 3 months ago

Answer "C" is correct as the question is not asking where its set, its asking what it protects or do.

upvoted 2 times

🗨️ **pnerella** 2 years, 4 months ago

it encrypts the database not the server. So answer is correct

upvoted 3 times

🗨️ **bfan95** 2 years, 4 months ago

Answer should be D

"or Azure SQL Database and Azure Synapse, the TDE protector is set at the server level and is inherited by all databases associated with that server."

Source: <https://docs.microsoft.com/en-us/azure/azure-sql/database/transparent-data-encryption-tde-overview?tabs=azure-portal>

upvoted 1 times

🗨️ **J\_ose** 1 year, 8 months ago

It can't encrypt the server but the data on the database

upvoted 1 times

🗨️ **nigho4real** 2 years, 5 months ago

I feel the answer is correct, according to this link <https://docs.microsoft.com/en-us/azure/azure-sql/database/transparent-data-encryption-tde-overview?tabs=azure-portal>

Transparent data encryption (TDE) helps protect Azure SQL Database, Azure SQL Managed Instance, and Azure Synapse Analytics against the threat of malicious offline activity by encrypting data at rest. It performs real-time encryption and decryption of the database, associated backups, and transaction log files at rest without requiring changes to the application.

it seems to deal with database at rest to me.

upvoted 1 times

🗨️ **azurearch** 2 years, 5 months ago

For Azure SQL Database and Azure Synapse, the TDE protector is set at the server level and is inherited by all databases associated with that server. , Answer : D

upvoted 1 times

🗨️ **Luis\_polanco** 2 years, 6 months ago

The documentation says: "It performs real-time encryption and decryption of the database,..." So C is correct.

upvoted 1 times

🗨️ **fluid\_clearncer** 2 years, 7 months ago

The answer is D

[https://docs.microsoft.com/en-us/sql/relational-databases/security/encryption/transparent-data-encryption?view=sql-server-ver15#:~:text=Transparent%20data%20encryption%20\(TDE\)%20encrypts,Designing%20a%20secure%20system.](https://docs.microsoft.com/en-us/sql/relational-databases/security/encryption/transparent-data-encryption?view=sql-server-ver15#:~:text=Transparent%20data%20encryption%20(TDE)%20encrypts,Designing%20a%20secure%20system.)

upvoted 2 times

You need to ensure that users use multi-factor authentication (MFA) when connecting to an Azure SQL database. Which type of authentication should you use?

- A. service principal authentication
- B. Azure Active Directory (Azure AD) authentication
- C. SQL authentication
- D. certificate authentication

**Suggested Answer:** B

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/authentication-mfa-ssms-overview>

Community vote distribution

B (100%)

 **Bolarin** Highly Voted 3 years, 5 months ago

Azure AD is correct  
upvoted 52 times

 **kiarash** Highly Voted 3 years, 1 month ago

B- AD , AZ900 is that you?  
upvoted 19 times

 **alluxenof** 2 years, 9 months ago

I think that AZ-900 is a good entry point for AI/DP -900 !  
upvoted 13 times

 **Bozho** 2 years, 5 months ago

Absolutely!  
<https://www.examttopics.com/exams/microsoft/dp-900/view/10/#>  
upvoted 2 times

 **AGTraining** Most Recent 6 months, 2 weeks ago

option B  
upvoted 1 times

 **chiiiweiii** 1 year, 6 months ago

Selected Answer: B

Correct Answer  
upvoted 2 times

 **Pinha** 1 year, 7 months ago

B. Other authentication types, such as service principal authentication, SQL authentication and certificate authentication, don't allow the use of MFA.  
upvoted 5 times

 **JA2018** 1 year, 11 months ago

Selected Answer: B

MFA required Azure AD. Covered in AZ900 exam  
upvoted 7 times

 **AndreLima** 1 year, 12 months ago

Make AZ-900 before makes DP-900. Its best.  
upvoted 3 times

 **haileysbeth** 2 years ago

Option B is right  
upvoted 1 times

 **Eltooth** 2 years, 2 months ago

**Selected Answer: B**

B is correct answer.

With CAP added on...

upvoted 1 times

  **RichMelbourne** 2 years, 5 months ago

Shown on test 15/04/2022

upvoted 5 times

  **SiDoCiOuS** 2 years, 5 months ago

**Selected Answer: B**

Appeared on the exam on 03/22/2022

Answer is correct!

upvoted 4 times

  **GuruKannan** 2 years, 7 months ago

AD is correct

upvoted 1 times

  **AzureDP900** 2 years, 11 months ago

Azure Active Directory (Azure AD) authentication perfect answer no second thought.

upvoted 3 times

  **GetulioJr** 3 years, 3 months ago

Not related to this exam

upvoted 1 times

  **mrjeet** 2 years, 8 months ago

does it matter? You should still know this info

upvoted 1 times

  **rogui99** 3 years, 6 months ago

MFA comes with Azure AD

upvoted 3 times

  **bobhichem** 3 years, 6 months ago

azure AD is the correct answer

upvoted 1 times

  **Vithal1979** 3 years, 11 months ago

azure AD comes with MFA

upvoted 9 times

What is a benefit of hosting a database on Azure SQL managed instance as compared to an Azure SQL database?

- A. built-in high availability
- B. native support for cross-database queries and transactions
- C. system-initiated automatic backups
- D. support for encryption at rest

**Suggested Answer:** B

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/features-comparison>

Community vote distribution

B (100%)

- 🗨️ **jasu** Highly Voted 4 years ago  
Reference document <https://docs.microsoft.com/en-us/azure/azure-sql/database/features-comparison>  
upvoted 45 times
- 🗨️ **cadio30** Highly Voted 3 years ago  
B is the correct answer "native support for cross-database queries and transactions"  
upvoted 25 times
- 🗨️ **AGTraining** Most Recent 6 months, 2 weeks ago  
option B  
upvoted 1 times
- 🗨️ **Ycombo** 9 months, 1 week ago  
How many pages or Qs should I atleast go thurgh to be confident that I have covered enough, I am at Page 17 atm ?  
upvoted 1 times
- 🗨️ **Anna\_Le** 9 months ago  
better go through all of the questions:)  
upvoted 3 times
- 🗨️ **Kans19** 1 year, 4 months ago  
Selected Answer: B  
ACD are available for both SQL db and SQL managed instance.  
While multiple writes and reads for SQL managed instance and only multi read alone is available for SQL db. So it's B  
upvoted 2 times
- 🗨️ **GaneshNats** 1 year, 4 months ago  
Selected Answer: B  
B is the right answer  
upvoted 3 times
- 🗨️ **chiiiweiii** 1 year, 6 months ago  
Selected Answer: B  
Correct Answer  
upvoted 2 times
- 🗨️ **Dusica** 1 year, 10 months ago  
B; other 3 are all in common for all SQL databases  
upvoted 5 times
- 🗨️ **JA2018** 1 year, 11 months ago  
Selected Answer: B  
Right response  
upvoted 3 times
- 🗨️ **SiDoCiOuS** 2 years, 5 months ago

**Selected Answer: B**

Appeared on the exam on 03/22/2022

Answer is correct!

upvoted 4 times

  **Koph** 2 years, 9 months ago

Tricky one but B is correct

upvoted 4 times

  **bluetail** 2 years, 9 months ago

No, it is D) support for encryption at rest. <https://docs.microsoft.com/en-us/azure/azure-sql/managed-instance/sql-managed-instance-paas-overview>

upvoted 1 times

  **iooj** 2 years, 8 months ago

No. TDE available in both. "B" is the correct answer. See here:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/features-comparison>

upvoted 4 times

  **Gbolatiti** 2 years, 11 months ago

I just passed this exam now...this question came out

upvoted 9 times

  **UniversalSoldier** 3 years, 1 month ago

Answer is correct, search cross-database on the below link

<https://docs.microsoft.com/en-us/azure/azure-sql/database/features-comparison>

upvoted 4 times

  **Adebowale** 3 years, 1 month ago

Good analysis and reference

upvoted 1 times

  **Crazy\_Angel** 3 years, 2 months ago

Cross-database/three-part name queries Yes, plus Elastic queries

Cross-database transactions Yes, within the instance.

Hence, I think B is the answer.

native support for cross-database queries and transactions

upvoted 1 times

  **razxssd** 3 years, 2 months ago

Ciao Ale, Iso, Ale

upvoted 2 times

HOTSPOT -

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

Hot Area:

### 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.	

Suggested Answer:

### 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.	

When you create a new server in Azure SQL Database or Azure Synapse Analytics named mysqlserver, for example, a server-level firewall blocks all access to the public endpoint for the server

Reference:

<https://docs.microsoft.com/en-us/azure/security/fundamentals/infrastructure-sql>

🗨️ **ExamPwnr** Highly Voted 4 years ago

Correct answer: "a server level firewall"

Azure SQL has it's own Firewall and firewall rules options

It is not "Azure Firewall" though:

"Azure Firewall is a managed, cloud-based network security service that protects your Azure Virtual Network resources. It's a fully stateful firewall as a service with built-in high availability and unrestricted cloud scalability."

upvoted 126 times

🗨️ **BrettusMaximus** 3 years, 11 months ago

It can't be server level as the database is clustered (multiple servers).

upvoted 4 times

🗨️ **jpgerek** 3 years, 8 months ago

<https://docs.microsoft.com/en-us/azure/security/fundamentals/infrastructure-sql#azure-sql-database-firewall>

upvoted 11 times

🗨️ **vince60370** 3 years, 6 months ago

Thanks for the link, the schema clearly states "server-level firewall (rules)"

upvoted 9 times

🗨️ **Jdmann** Highly Voted 4 years ago

Correct Answer is B

<https://docs.microsoft.com/en-us/azure/azure-sql/database/firewall-create-server-level-portal-quickstart#:~:text=A%20server%2Dlevel%20IP%20firewall%20rule%20can%20open%20port%201433,a%20range%20of%20IP%20addresses.&text=By%20de>

upvoted 19 times

🗨️ **AGTraining** Most Recent 6 months, 2 weeks ago

option B

upvoted 1 times

🗨️ **Tnob92** 8 months ago

Answer is B. Database-level rules are not configured by default and need to be set up for more granular control.

upvoted 1 times

🗨️ **Suraj2801** 1 year, 5 months ago

CORRECT ANSWER...

upvoted 3 times

🗨️ **chiiiweiii** 1 year, 6 months ago

Correct Answer

upvoted 2 times

🗨️ **dk87** 2 years, 4 months ago

Exam: 05/18/2022

upvoted 6 times

🗨️ **nicolase** 2 years, 6 months ago

Select Set server firewall on the toolbar. The Firewall settings page for the server opens.

upvoted 2 times

🗨️ **rick245** 2 years, 8 months ago

Can someone please clarify this aspect regarding Azure SQL Database?

From Microsoft website I can see two different graph of how firewall works:

1) <https://docs.microsoft.com/en-us/azure/security/fundamentals/infrastructure-sql>

Internet > Server-Level Firewall rules > Database-Level Firewall rules > Database access

2) <https://docs.microsoft.com/en-us/azure/azure-sql/database/firewall-configure>

Internet > Database-Level Firewall rules > Server-Level Firewall rules > Database access

By looking at the different algorithms, if we have an IP allowed in both Database-Level Firewall (for Database1) and Server-Level Firewall (which allows access to every database) we have two different scenario

if 1) is the correct implementation, we can access every database (Database1, Database2, Database3, Database4)

if 2) is the correct implementation, we can access only Database1

Which is the correct one?

upvoted 2 times

🗨️ **osaid** 2 years, 8 months ago

Correct answer is B. Reference is given below:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/firewall-configure>

When you create a new server in Azure SQL Database or Azure Synapse Analytics named mysqlserver, for example, a server-level firewall blocks all access to the public endpoint for the server (which is accessible at mysqlserver.database.windows.net).

upvoted 3 times

🗨️ **Koph** 2 years, 9 months ago

Server level firewall

upvoted 1 times

🗨️ **cadio30** 3 years ago

Server Firewall is the correct answer but one can enhance the security by executing a script to add a database firewall

upvoted 2 times

🗨️ **priy\_0987** 3 years, 1 month ago

itt should be B

upvoted 1 times

🗨️ **PhillipH** 3 years, 2 months ago

azure firewall - by default

upvoted 2 times

🗨️ **robinho0129** 3 years, 2 months ago

Por defecto es Azure Firewall, el Server firewall debe configurarse, por ende la correcta es la Azuere Firewall.

upvoted 1 times

  **danielloesc** 3 years, 3 months ago

correct

upvoted 2 times

  **saadnadir** 3 years, 3 months ago

"each" database = > database firewall

upvoted 2 times

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

**Suggested Answer:** A

Reference:

<https://docs.microsoft.com/en-us/sql/ssdt/project-oriented-offline-database-development?view=sql-server-ver15>

Community vote distribution

A (100%)

 **learner301** Highly Voted 4 years ago

I agree with answer and link  
upvoted 56 times

 **Aneta\_ane** Highly Voted 2 years, 9 months ago

This question came up on the exam today  
upvoted 19 times

 **D\_Mo\_ZY** 2 years, 6 months ago

Was in the exam (18/03/2022)  
upvoted 15 times

 **AGTraining** Most Recent 6 months, 2 weeks ago

option A  
upvoted 1 times

 **chiiiweiii** 1 year, 6 months ago

Selected Answer: A

Correct Answer  
upvoted 6 times

 **Joanna0** 1 year, 7 months ago

Selected Answer: A

<https://learn.microsoft.com/en-us/sql/ssdt/project-oriented-offline-database-development?view=sql-server-ver16>  
upvoted 2 times

 **Teamdev** 2 years, 5 months ago

Agree, correct answer  
upvoted 3 times

 **nicolase** 2 years, 7 months ago

Selected Answer: A

SSDT, you can create an offline database project and implement schema changes by adding, modifying or deleting the definitions of objects (represented by scripts) in the project  
upvoted 10 times

 **JA2018** 1 year, 11 months ago

it is confirmed by the following URL:

<https://learn.microsoft.com/en-us/sql/ssdt/sql-server-data-tools?view=sql-server-ver16>

upvoted 1 times

 **theLazyGuy** 2 years, 10 months ago

...Using SSDT, you can create an offline database project...

(<https://docs.microsoft.com/en-us/sql/ssdt/project-oriented-offline-database-development?view=sql-server-ver15>)

upvoted 6 times

  **cadio30** 3 years ago

SSDT can create SSRS, SSIS, SSAS and database development by using the "database project".

upvoted 6 times

  **p6wruv** 3 years ago

Both A and B can do this Offline. Install the Management Studio and you will see that you can use visual tools and work offline.

upvoted 6 times

  **reeftu** 3 years, 2 months ago

I agree with the answer. The key word is OFFLINE.

upvoted 4 times

  **Armandoo** 3 years, 3 months ago

Correct Answer: SSDT

SQL Server Data Tools (SSDT) is a modern development tool for building SQL Server relational databases, databases in Azure SQL, Analysis Services (AS) data models, Integration Services (IS) packages, and Reporting Services (RS) reports. With SSDT, you can design and deploy any SQL Server content type with the same ease as you would develop an application in Visual Studio.

upvoted 5 times

  **Shiv948** 3 years, 3 months ago

Answer A is correct.

- SSDT is for designing stuff related to Integration Services, Analysis Services, and Reporting Services. You can also use it to design and deploy databases to SQL Server instances.

SSMS is a way to manage SQL server instances.

upvoted 4 times

  **ramjisriram** 3 years, 6 months ago

Given answer is correct.

upvoted 5 times

  **mandyfervor** 3 years, 7 months ago

I agree with the answer A as well.

upvoted 3 times

  **gbaren** 3 years, 7 months ago

Weird question. SSDT is not for database development, it includes SSIS for integration development.

upvoted 3 times

  **korapavan433** 3 years, 6 months ago

no that's wrong, SSDT not only for ssis. its also for SSAS etc..in SSAS you can design the model

upvoted 2 times

  **crazycga** 3 years, 6 months ago

I'd honestly never thought of using SSDT in this manner. I usually only use SSIS and SSRS...

upvoted 1 times

  **ramjisriram** 3 years, 6 months ago

That's fine if you don't use the evolved SSDT, as per the documentation SSDT can be used.

upvoted 1 times

  **rmn900** 3 years, 10 months ago

A

[https://docs.microsoft.com/en-us/sql/ssdt/download-sql-server-data-tools-ssdt?view=sql-server-ver15#:~:text=SQL%20Server%20Data%20Tools%20\(SSDT\)%20is%20a%20modern%20development%20tool,Reporting%20Services%20\(RS\)%20reports.](https://docs.microsoft.com/en-us/sql/ssdt/download-sql-server-data-tools-ssdt?view=sql-server-ver15#:~:text=SQL%20Server%20Data%20Tools%20(SSDT)%20is%20a%20modern%20development%20tool,Reporting%20Services%20(RS)%20reports.)

upvoted 2 times

DRAG DROP -

Match the security components to the appropriate scenarios.

To answer, drag the appropriate component from the column on the left to its scenario on the right. Each component may be used once, more than once, or not at all.

NOTE: Each correct match is worth one point.

Select and Place:

### Components

Authentication

Firewall

Encryption

### Answer Area

Prevent access to an Azure SQL database from another network.

Support Azure Active Directory (Azure AD) sign-ins to an Azure SQL database.

Ensure that sensitive data never appears as plain text in an Azure SQL database.

### Suggested Answer:

#### Components

Authentication

Firewall

Encryption

#### Answer Area

Firewall

Prevent access to an Azure SQL database from another network.

Authentication

Support Azure Active Directory (Azure AD) sign-ins to an Azure SQL database.

Encryption

Ensure that sensitive data never appears as plain text in an Azure SQL database.

Reference:

<https://docs.microsoft.com/en-us/sql/database-engine/configure-windows/configure-a-windows-firewall-for-database-engine-access?view=sql-server-ver15> <https://docs.microsoft.com/en-us/azure/azure-sql/database/authentication-aad-overview>  
<https://docs.microsoft.com/en-us/azure/azure-sql/database/always-encrypted-certificate-store-configure>

 **learner301** Highly Voted 4 years ago

I agree with answers

upvoted 76 times

 **agautam** Highly Voted 3 years, 11 months ago

Agreed

upvoted 18 times

 **AGTraining** Most Recent 6 months, 2 weeks ago

Firewall,

Authentication,

Encryption

upvoted 1 times

 **Ejake** 1 year, 5 months ago

Answer is correct

upvoted 2 times

 **chiiiweiii** 1 year, 6 months ago

Correct Answer

upvoted 2 times

 **Asma1990** 1 year, 11 months ago

absolutely correct!

upvoted 4 times

 **roshanazMaster** 2 years, 4 months ago

correct

upvoted 2 times

 **Teamdev** 2 years, 5 months ago

correct answers

upvoted 1 times

🗨️ 👤 **SiDoCiOuS** 2 years, 5 months ago  
Appeared on the exam on 03/22/2022  
Answers are correct! Firewall, Authentication, Encryption  
upvoted 3 times

🗨️ 👤 **Amadh** 2 years, 6 months ago  
Agree with the answer  
upvoted 2 times

🗨️ 👤 **Koph** 2 years, 9 months ago  
Correct  
upvoted 1 times

🗨️ 👤 **martin898** 2 years, 11 months ago  
Agreed  
upvoted 1 times

🗨️ 👤 **Terminator10** 3 years ago  
Agreed  
upvoted 3 times

🗨️ 👤 **PankajBadhiye** 3 years, 1 month ago  
Agreed  
upvoted 1 times

🗨️ 👤 **Adechrist** 3 years, 1 month ago  
Agreed  
upvoted 2 times

🗨️ 👤 **markwu01** 3 years, 2 months ago  
Agreed  
upvoted 3 times

🗨️ 👤 **Pranav\_K** 3 years, 2 months ago  
Agreed  
upvoted 3 times

You have a transactional application that stores data in an Azure SQL managed instance.  
When should you implement a read-only database replica?

- A. You need to generate reports without affecting the transactional workload.
- B. You need to audit the transactional application.
- C. You need to implement high availability in the event of a regional outage.
- D. You need to improve the recovery point objective (RPO).

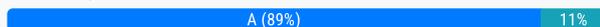
**Suggested Answer: A**

Use read-only replicas to offload read-only query workloads.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/read-scale-out>

Community vote distribution



**ExamPwnr** Highly Voted 4 years ago

Why C? It is not High Availability if you cannot write to the database. If your primary database goes down, HA means that users would continue to work as normal, but they cannot commit changes in this case (Read-Only). Answer should be only: A.  
upvoted 127 times

**SQLHell** 2 years, 9 months ago

All the other options are managed by Azure SQL database  
upvoted 3 times

**learner301** Highly Voted 4 years ago

I would go for A, read only replica is created to offload part of transactions (backup, reporting...) in order to free up resources on DB.  
If you want to achieve High availability use Geo redundancy  
upvoted 70 times

**fdlsys** 3 years, 8 months ago

Secondary instance in HA is read-only for users but write-able for DB sync.  
So, answer is "C" - that is the actual purpose of the secondary replica(s).

Referenced URL merely suggests that replicas can and should be used to offload heavy read-only (reports) workload from the primary instance, but that does not define the original purpose.  
upvoted 3 times

**jimakana** 3 years, 7 months ago

How does this make sense to the "in the event of a regional outage" part in the answer?  
upvoted 2 times

**AGTraining** Most Recent 6 months, 2 weeks ago

option A  
upvoted 1 times

**ItzFede** 1 year, 2 months ago

The correct answer is A. You should implement a read-only database replica when you need to generate reports without affecting the transactional workload.

By creating a read-only replica of the transactional database, you can offload the reporting workload to the replica, ensuring that the performance of the transactional workload is not impacted. This allows you to generate reports and perform analytical queries on the replica without affecting the availability or responsiveness of the primary transactional database.  
upvoted 4 times

**chiiiweiii** 1 year, 6 months ago

Selected Answer: A  
Correct Answer

upvoted 2 times

🗨️ **milan92stankovic** 1 year, 7 months ago

**Selected Answer: A**

Clear as a tear - A!

upvoted 2 times

🗨️ **Stuurwiel1** 1 year, 7 months ago

**Selected Answer: A**

You should implement a read-only database replica when you need to generate reports without affecting the transactional workload.

A read-only replica is a copy of the database that can be used for reporting or other read-only workloads without affecting the performance of the transactional application. This allows you to offload reporting and other read-only workloads to the replica, which can improve the performance of the transactional application.

upvoted 6 times

🗨️ **King\_Lam** 1 year, 7 months ago

A is correct.

upvoted 2 times

🗨️ **thakurvi** 1 year, 8 months ago

**Selected Answer: A**

Its A for sure . C is debatable if kept for active passive config as DR

upvoted 2 times

🗨️ **Faize** 1 year, 9 months ago

The right Answer is A.

upvoted 2 times

🗨️ **tizona** 1 year, 9 months ago

**Selected Answer: A**

right answer is A

upvoted 1 times

🗨️ **WWDITDA** 2 years, 2 months ago

**Selected Answer: A**

Absolutely correct.

upvoted 3 times

🗨️ **cmohl2013** 2 years, 3 months ago

**Selected Answer: A**

C ist not correct: operational DBs must provide write access to be fully functional, a read-only DB will not cover for that and availability for writes is not maintained.

upvoted 4 times

🗨️ **DEMOLKID** 2 years, 3 months ago

conflicto de interes

upvoted 1 times

🗨️ **bfan95** 2 years, 4 months ago

**Selected Answer: A**

Answer: A

upvoted 2 times

🗨️ **RichMelbourne** 2 years, 5 months ago

Shown on test 15/04/2022

upvoted 3 times

🗨️ **medinacau** 2 years, 5 months ago

**Selected Answer: A**

A is correct

upvoted 1 times

HOTSPOT -

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

Hot Area:

### 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.	

Suggested Answer:

### 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.	

ANath Highly Voted 3 years, 5 months ago

correct

upvoted 45 times

jasu Highly Voted 4 years ago

Reference Reading: <https://docs.microsoft.com/en-us/sql/relational-databases/indexes/clustered-and-nonclustered-indexes-described?view=sql-server-ver15#:~:text=Clustered%20indexes%20sort%20and%20store,stored%20in%20only%20one%20order.>

upvoted 19 times

AGTraining Most Recent 6 months, 2 weeks ago

option D

upvoted 2 times

globy118 1 year, 2 months ago

Was on exam 24/6/2023

upvoted 2 times

Ejake 1 year, 5 months ago

Correct

upvoted 2 times

chiiweiii 1 year, 6 months ago

Correct Answer

upvoted 2 times

SAIVEENA 1 year, 9 months ago

4th Option is correct.

1.Relational databases are static schema so option A is wrong.

2.data stored as key/value pairs is JSON (semi-structured data)relationalDB are structured data. Option B is wrong

3. Storing large images and videos is unstructured data.Option c is wrong.

upvoted 12 times

ria97 2 years, 1 month ago

Was in exam 11/08. Correct answer

upvoted 4 times

Breeze21 2 years, 2 months ago

Exam: 19.07.22

upvoted 3 times

🗨️ 👤 **passenger** 2 years, 3 months ago

For this question I would select the 4th option however, I would only do that because this is the most correct out of all offered. The correct answer here is in fact incorrect because you do not have to resort to a relational database if you want to enforce ACID. There are many tools that can enforce ACID in other NoSQL databases, it does not makes relational databases the only way.

upvoted 3 times

🗨️ 👤 **print>HelloWorld** 2 years, 4 months ago

Correct

upvoted 2 times

🗨️ 👤 **roshanazMaster** 2 years, 4 months ago

ACID so correct

upvoted 2 times

🗨️ 👤 **RichMelbourne** 2 years, 5 months ago

Shown on test 15/04/2022

upvoted 6 times

🗨️ 👤 **Teamdev** 2 years, 5 months ago

It's correct !!

upvoted 2 times

🗨️ 👤 **SiDoCiOuS** 2 years, 5 months ago

Appeared on the exam on 03/22/2022

Answer is correct! Strong Consistency

upvoted 3 times

🗨️ 👤 **RajeshAzure** 2 years, 6 months ago

ACID rules, where C is for consistency

upvoted 3 times

🗨️ 👤 **WigglyPork** 2 years, 7 months ago

correct, basic things in a relational database

upvoted 1 times

HOTSPOT -

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

Hot Area:

### Answer Area

Relational data uses  to enforce relationships between different tables.

collections
columns
keys
partitions

Suggested Answer:

### Answer Area

Relational data uses  to enforce relationships between different tables.

collections
columns
keys
partitions

Reference:

<https://teachcomputerscience.com/relational-databases/>

- 👤 **Vithal1979** Highly Voted 3 years, 11 months ago  
 Again RDBMS concepts primary key, unique key & foreign key concepts  
 upvoted 49 times
- 👤 **olowoyinka** Highly Voted 3 years, 5 months ago  
 correct answer  
 upvoted 24 times
- 👤 **AGTraining** Most Recent 6 months, 2 weeks ago  
 keys corect  
 upvoted 1 times
- 👤 **globy118** 1 year, 2 months ago  
 Was on exam 24/6/2023  
 upvoted 2 times
- 👤 **Ejake** 1 year, 5 months ago  
 correct  
 upvoted 3 times
- 👤 **chiiiweiii** 1 year, 6 months ago  
 Correct Answer  
 upvoted 2 times
- 👤 **dk87** 2 years, 4 months ago  
 Exam: 05/18/2022  
 upvoted 6 times
- 👤 **roshanazMaster** 2 years, 4 months ago  
 Keys correct answer  
 upvoted 4 times
- 👤 **Teamdev** 2 years, 5 months ago  
 correct !  
 upvoted 2 times
- 👤 **Koph** 2 years, 9 months ago

Keys is correct  
upvoted 3 times

🗨️ **ashb20** 2 years, 10 months ago  
correct  
upvoted 4 times

🗨️ **Gbolatiti** 2 years, 11 months ago  
is it that no one has written this exam recently?  
upvoted 1 times

🗨️ **JpLZa** 1 year, 6 months ago  
Why you ask? Answer is valid  
upvoted 1 times

🗨️ **MFO\_FM** 3 years ago  
Correct Answer: Keys  
No arguments  
upvoted 2 times

🗨️ **Armandoo** 3 years, 3 months ago  
Correct Answer: Keys  
relational databases use PK and FK to enforce relationships between tables.  
upvoted 3 times

🗨️ **catalene** 3 years, 4 months ago  
Why not "Columns"? You use columns to build a relationship and then the system create a "foreign key". Key definition is other..  
upvoted 2 times

🗨️ **ndatta** 3 years, 1 month ago  
because we dont use any random columns to create a relationship. we create entity-relation diagram to enforce the relationship among tables and that is done based on the primary key-foreign key relationships  
upvoted 3 times

🗨️ **ANath** 3 years, 5 months ago  
Correct  
upvoted 1 times

🗨️ **nkav** 3 years, 6 months ago  
correct answer. keys!  
upvoted 3 times

HOTSPOT -

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

Hot Area:

### Answer Area

▼	is a virtual table that contains content defined by a query.
A heap	
A stored procedure	
A view	
An index	

### Answer Area

Suggested Answer:

▼	is a virtual table that contains content defined by a query.
A heap	
A stored procedure	
A view	
An index	

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/views/views>

🗨️ **shinz** Highly Voted 3 years, 6 months ago

heap: tables: without a clustered index.

Stored procedure: a sub-routine available to access application in RDBMS

A view: is a table that is being created based on a user query

An index: a data structure that improves the speed of data retrieval operations

The answer is obvious, A View

upvoted 50 times

🗨️ **kaleidoscope** Highly Voted 3 years, 3 months ago

Correct Answer

upvoted 12 times

🗨️ **AGTraining** Most Recent 6 months, 2 weeks ago

view correct

upvoted 2 times

🗨️ **chiiiweiii** 1 year, 6 months ago

Correct Answer

upvoted 2 times

🗨️ **geeeogeee** 2 years, 3 months ago

Yes answer is correct

upvoted 3 times

🗨️ **dk87** 2 years, 4 months ago

Exam: 05/18/2022

upvoted 3 times

🗨️ **Teamdev** 2 years, 5 months ago

correct answer

upvoted 2 times

🗨️ **Ravikiransattineni** 2 years, 8 months ago

Answer is Correct

upvoted 2 times

🗨️ 👤 **Koph** 2 years, 9 months ago

A view is correct

upvoted 1 times

🗨️ 👤 **MFO\_FM** 3 years ago

Correct

upvoted 3 times

🗨️ 👤 **ANath** 3 years, 5 months ago

correct

upvoted 4 times

🗨️ 👤 **tochy** 3 years, 5 months ago

View is correct

upvoted 3 times

🗨️ 👤 **overhill** 3 years, 5 months ago

not comfortable with the virtual table naming but okay

upvoted 1 times

🗨️ 👤 **Raghu22** 3 years, 5 months ago

Correct Answer

upvoted 2 times

🗨️ 👤 **Andres91salazar** 3 years, 7 months ago

A view is the correct answer

upvoted 3 times

🗨️ 👤 **GHNZED** 3 years, 8 months ago

the answer is Views

upvoted 4 times

🗨️ 👤 **teweve19** 3 years, 9 months ago

Views is the answer

upvoted 3 times

You need to query a table named Products in an Azure SQL database.

Which three requirements must be met to query the table from the internet? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. You must be assigned the Reader role for the resource group that contains the database.
- B. You must have SELECT access to the Products table.
- C. You must have a user in the database.
- D. You must be assigned the Contributor role for the resource group that contains the database.
- E. Your IP address must be allowed to connect to the database.

**Suggested Answer:** BCE

Incorrect Answers:

A, D: Resource group permissions is not required to query an Azure SQL database table.

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/security/authentication-access/getting-started-with-database-engine-permissions?view=sql-server-ver15>

Community vote distribution

BCE (100%)

 **alonso3** Highly Voted 2 years, 4 months ago

BCE is correct

upvoted 26 times

 **fusioner** Highly Voted 2 years ago

Selected Answer: BCE

BCE is Correct Answer

upvoted 9 times

 **AGTraining** Most Recent 6 months, 2 weeks ago

BCE is correct

upvoted 1 times

 **AZFabio** 1 year ago

Selected Answer: BCE

BCE: your IP must be allowed in DB firewall and then you need a DB user with select privs on the disired table.

upvoted 2 times

 **poli123** 1 year, 1 month ago

Selected Answer: BCE

B,C,E are correct

upvoted 1 times

 **akhil5432** 1 year, 1 month ago

Selected Answer: BCE

All 3 options are correct

upvoted 1 times

 **KM9** 1 year, 1 month ago

Selected Answer: BCE

correct

upvoted 1 times

 **Prajai** 1 year, 2 months ago

if you have trouble remembering

key words: "must have" , "must have" , and "IP"

upvoted 4 times

🗨️ 👤 **chiiweiii** 1 year, 6 months ago

Selected Answer: BCE

Correct Answers

upvoted 3 times

🗨️ 👤 **opek** 1 year, 10 months ago

Selected Answer: BCE

Correct

upvoted 6 times

🗨️ 👤 **Emma\_123** 2 years, 3 months ago

BCE is correct

upvoted 5 times

🗨️ 👤 **NOLUPIKOLI** 2 years, 3 months ago

CORRECT

upvoted 3 times

🗨️ 👤 **bfan95** 2 years, 4 months ago

BCE is correct

upvoted 4 times

Which clause should you use in a SELECT statement to combine rows in one table with rows in another table?

- A. SET
- B. VALUES
- C. KEY
- D. JOIN

**Suggested Answer:** D

Reference:

<https://www.tutorialspoint.com/sql/sql-using-joins.htm>

Community vote distribution

D (100%)

🗳️ **svetlanam** Highly Voted 2 years, 4 months ago

Correct.

upvoted 18 times

🗳️ **Manoj6** Highly Voted 2 years, 3 months ago

Join is Correct

upvoted 8 times

🗳️ **AGTraining** Most Recent 6 months, 2 weeks ago

option D

upvoted 1 times

🗳️ **AZFabio** 1 year ago

Selected Answer: D

join - correct

upvoted 1 times

🗳️ **Murtuza** 1 year, 5 months ago

key word combine = join

upvoted 2 times

🗳️ **chiiiweiii** 1 year, 6 months ago

Selected Answer: D

Correct Answer

upvoted 2 times

🗳️ **Hari46** 1 year, 7 months ago

Selected Answer: D

JOIN is correct

upvoted 2 times

🗳️ **JA2018** 1 year, 11 months ago

Selected Answer: D

Agreed, please refer to following:

<https://learn.microsoft.com/en-us/sql/relational-databases/performance/joins?view=sql-server-ver16>

upvoted 2 times

🗳️ **Dasan** 2 years ago

That is correct

upvoted 2 times

🗳️ **Nivep** 2 years, 4 months ago

Selected Answer: D

Correct

upvoted 3 times

Your company needs to ensure that an Azure virtual machine can connect to Azure SQL databases without exposing the databases to the internet.

What should you use?

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

**Suggested Answer: D**

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

Reference:

<https://techcommunity.microsoft.com/t5/azure-database-support-blog/azure-sql-db-private-link-private-endpoint-connectivity/ba-p/1235573>

*Community vote distribution*

D (100%)

 **bfan95** Highly Voted 2 years, 4 months ago

**Selected Answer: D**

Answer: D

upvoted 13 times

 **jellybiscuit** Highly Voted 2 years ago

**Selected Answer: D**

Azure Private Link

upvoted 7 times

 **AGTraining** Most Recent 6 months, 2 weeks ago

option D

upvoted 1 times

 **Yomzie** 6 months, 3 weeks ago

Azure Private Link (Endpoint)

upvoted 2 times

 **chiiiweiii** 1 year, 6 months ago

**Selected Answer: D**

Correct Answer

upvoted 3 times

 **JA2018** 1 year, 11 months ago

**Selected Answer: D**

<https://learn.microsoft.com/en-us/azure/azure-sql/database/private-endpoint-overview?view=azuresql>

upvoted 3 times

 **bfan95** 2 years, 4 months ago

D is correct

upvoted 4 times

 **svetlanam** 2 years, 4 months ago

Correct.

upvoted 4 times

DRAG DROP -

Match the Azure SQL services to the appropriate use cases.

To answer, drag the appropriate service from the column on the left to its use case on the right. Each service may be used once, more than once, or not at all.

NOTE: Each correct match is worth one point.

Select and Place:

**Azure SQL services**    **Answer Area**

Azure SQL Database elastic pool		A serverless configuration
Azure SQL Database single database		A database that has the highest compatibility with on-premises Microsoft SQL Server
Azure SQL Managed instance		A low-cost, low-maintenance database

**Azure SQL services**    **Answer Area**

**Suggested Answer:**

Azure SQL Database elastic pool	Azure SQL Database single database	A serverless configuration
Azure SQL Database single database	Azure SQL Managed instance	A database that has the highest compatibility with on-premises Microsoft SQL Server
Azure SQL Managed instance	Azure SQL Database elastic pool	A low-cost, low-maintenance database

Box 1: Azure SQL Database single database

Serverless is a compute tier for single databases in Azure SQL Database that automatically scales compute based on workload demand and bills for the amount of compute used per second. The serverless compute tier also automatically pauses databases during inactive periods when only storage is billed and automatically resumes databases when activity returns.

Scenarios well suited for serverless compute

Single databases with intermittent, unpredictable usage patterns interspersed with periods of inactivity, and lower average compute utilization over time.

Single databases in the provisioned compute tier that are frequently rescaled and customers who prefer to delegate compute rescaling to the service.

New single databases without usage history where compute sizing is difficult or not possible to estimate prior to deployment in SQL Database.

Box 2: Azure SQL Managed Instance

Azure SQL Managed Instance is the intelligent, scalable cloud database service that combines the broadest SQL Server database engine compatibility with all the benefits of a fully managed and evergreen platform as a service.

Box 3: Azure SQL Database elastic pool

Azure SQL Database elastic pools are a simple, cost-effective solution for managing and scaling multiple databases that have varying and unpredictable usage demands. The databases in an elastic pool are on a single server and share a set number of resources at a set price. Elastic pools in SQL Database enable software as a service (SaaS) developers to optimize the price performance for a group of databases within a prescribed budget while delivering performance elasticity for each database.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/serverless-tier-overview?view=azuresql> <https://docs.microsoft.com/en-us/azure/azure-sql/database/elastic-pool-overview?view=azuresql> <https://docs.microsoft.com/en-us/azure/azure-sql/managed-instance/sql-managed-instance-paas-overview>

 **brainiacATH** Highly Voted 1 year, 12 months ago

Correct answer should be single, MI, single. Elastic Pool uses a server, and is thus not serverless, and the third point refers to 'a database' while Elastic Pool is for scaling multiple databases.

upvoted 16 times

🗨️ 👤 **okayhey** 6 months, 1 week ago

All azure services use a server... Lol it just depends on the type of cloud model what your responsibilities are. Elastic pool uses shared resources on a server, but you don't manage the server itself, nor the OS. You do manage scaling, which makes it PaaS and not Serverless. Single database can dynamically scale up and down, making it a Serverless solution.

Single, MI, single use correct, but not for the reasons you gave

upvoted 1 times

🗨️ 👤 **AnonymousJhb** 1 year, 11 months ago

Single, MI, single.

First option is single because MS:

Serverless is a compute tier for single databases in Azure SQL Database that automatically scales compute based on workload demand and bills for the amount of compute used per second. The serverless compute tier also automatically pauses databases during inactive periods when only storage is billed and automatically resumes databases when activity returns.

upvoted 3 times

🗨️ 👤 **j888** 1 year, 10 months ago

You can choose serverless or provisioned resources when creating the single database in Azure SQL.

<https://learn.microsoft.com/en-us/azure/azure-sql/database/single-database-create-quickstart?view=azuresql&tabs=azure-portal>.

"Each has its own service tier within the DTU-based purchasing model or vCore-based purchasing model and a compute size defining the resources allocated to the database engine."

<https://learn.microsoft.com/en-us/azure/azure-sql/database/single-database-overview?view=azuresql>

So I doubt this fall under serverless.

Elastic on the other hand is pure serverless, and the main purpose of Elastic pool is the ability to scale when needed which in turn will save cost. So I will choose Elastic, MI, Elastic.

upvoted 3 times

🗨️ 👤 **j888** 1 year, 10 months ago

Actually, the given answer is correct.

This statement did say single database as serverless.

<https://learn.microsoft.com/en-us/azure/azure-sql/database/serverless-tier-overview?view=azuresql>

upvoted 1 times

🗨️ 👤 **SwissAir** Highly Voted 2 years ago

Shouldn't it be Elastic Pool, MI then Single Db?

upvoted 8 times

🗨️ 👤 **Sr18** 4 months, 2 weeks ago

No, Elastic Pool is not serverless. So for first one it will be Single.

Second one is MI.

Where as 3rd one should be MI because we are talking about single database not databases. If you will have multiple databases then mostly Elastic Pool will be cheaper as all will share same resources.

upvoted 1 times

🗨️ 👤 **Sr18** 4 months, 2 weeks ago

My Bad Typo!!!,

Correct Ans: Single, MI, Single

upvoted 1 times

🗨️ 👤 **virtualcharm** 1 year, 5 months ago

Yes, EP, MI and single. The last two options want a database answer. The first options wants a configuration answer, elastic pool is a config.

upvoted 4 times

🗨️ 👤 **just\_karlijn** 1 year, 11 months ago

Not a pro, but I had the same idea.

upvoted 1 times

🗨️ 👤 **JA2018** 1 year, 11 months ago

For Elastic pool, all the databases are hosted on a single server to allow sharing of system resources (e.g. CPU, RAM, disk, etc)

upvoted 1 times

🗨️ 👤 **AGTraining** Most Recent 6 months, 2 weeks ago

EP,  
MI  
single  
upvoted 1 times

🗨️ 👤 **m0j7zdg** 7 months, 2 weeks ago

Azure SQL DB Single Instance - serverless

<https://learn.microsoft.com/en-us/azure/azure-sql/database/serverless-tier-overview?view=azuresql&tabs=general-purpose>

Serverless is a compute tier for single databases in Azure SQL Database that automatically scales compute based on workload demand and bills for the amount of compute used per second.

Azure SQL Managed Instance - highest compatible with on-premise SQL server (Actually it's SQL server on VM but that's not an option here)

<https://learn.microsoft.com/en-us/azure/azure-sql/managed-instance/sql-managed-instance-paas-overview?view=azuresql>

[https://azure.microsoft.com/en-us/products/virtual-machines/sql-](https://azure.microsoft.com/en-us/products/virtual-machines/sql-server#:~:text=SQL%20Server%20on%20Azure%20Virtual%20Machines%20is%20100%20percent%20compatible,Pro%20Linux%2C%20and%20Windows%20)

[server#:~:text=SQL%20Server%20on%20Azure%20Virtual%20Machines%20is%20100%20percent%20compatible,Pro%20Linux%2C%20and%20Windows%20](https://azure.microsoft.com/en-us/products/virtual-machines/sql-server#:~:text=SQL%20Server%20on%20Azure%20Virtual%20Machines%20is%20100%20percent%20compatible,Pro%20Linux%2C%20and%20Windows%20)

low cost low maint DB - Azure SQL DB Elastic Pool

upvoted 2 times

🗨️ 👤 **Kemsmems** 10 months, 1 week ago

"Serverless is a compute tier for single databases in Azure SQL Database"

<https://learn.microsoft.com/en-us/azure/azure-sql/database/serverless-tier-overview?view=azuresql&tabs=general-purpose>.

upvoted 3 times

🗨️ 👤 **Pixan** 10 months, 2 weeks ago

Hi Everyone!!

Join ET and get actual and valid study material: <https://examsttopics.quora.com/> and pass your exam in first attempt. Study Smart Not Hard

upvoted 1 times

🗨️ 👤 **aadil111** 1 year, 5 months ago

1. Azure SQL Database Elastic Pool: C. A low-cost, low-maintenance database.

- An Azure SQL Database Elastic Pool allows you to manage and scale multiple databases with a shared set of resources. It is a cost-effective option for managing databases with varying and unpredictable usage patterns.

2. Azure SQL Database Single Database: B. A database that has the highest compatibility with on-premises Microsoft SQL Server.

- Azure SQL Database Single Database is a fully managed, highly available, and scalable relational database service. It provides the highest compatibility with on-premises Microsoft SQL Server and supports most of the features available in SQL Server.

3. Azure SQL Managed Instance: A. A serverless config.

- Azure SQL Managed Instance is a fully managed platform-as-a-service (PaaS) solution that provides SQL Server compatibility. It offers a serverless configuration that automatically scales compute and storage resources based on the workload demands.

upvoted 5 times

🗨️ 👤 **barcodepiglet** 10 months, 2 weeks ago

I don't think is correct. Azure SQL Managed Instance has the highest compability with on-premises MSFT SQL Server.

upvoted 4 times

🗨️ 👤 **chiiiweiii** 1 year, 6 months ago

Correct Answer

upvoted 2 times

🗨️ 👤 **ravick4u** 1 year, 7 months ago

The answer is correct (Key point is "serverless" which is for Azure SQL Database single database and "low cost" which is for elastic pool). I can not update the image here but when you try to create new Azure SQL database read the description on the portal.

<https://learn.microsoft.com/en-us/azure/azure-sql/database/serverless-tier-overview?view=azuresql>

<https://learn.microsoft.com/en-us/azure/azure-sql/database/elastic-pool-overview?view=azuresql>

upvoted 3 times

🗨️ 👤 **Dusica** 1 year, 9 months ago

correct: Single, Managed, Elastic

upvoted 2 times

🗨️ 👤 **Maddy** 1 year, 10 months ago

Single,MI,Single

upvoted 2 times

🗨️ 👤 **OldSchool** 1 year, 11 months ago

1. Single Instance "Serverless is a compute tier for single databases in Azure SQL Database..." <https://learn.microsoft.com/en-us/azure/azure-sql/database/serverless-tier-overview?view=azuresql>

2. Managed Instance

3. Elastic Pool "Azure SQL Database elastic pools are a simple, cost-effective solution for managing and scaling multiple databases that have varying and unpredictable usage demands." <https://learn.microsoft.com/en-us/azure/azure-sql/database/elastic-pool-overview?view=azuresql>

upvoted 5 times

🗨️ 👤 **PatYeo** 1 year, 11 months ago

The model answer appears to be correct based on the supporting statements and links :

1) Single Instance

2) Managed Instance

3) Elastic Pool

upvoted 7 times

🗨️ 👤 **fusioner** 2 years ago

Correct Answer

upvoted 2 times

Which T-SQL statement should be used to instruct a database management system to use an index instead of performing a full table scan?

- A. SELECT
- B. WHERE
- C. JOIN

**Suggested Answer: A**

Table hints override the default behavior of the Query Optimizer for the duration of the data manipulation language (DML) statement by specifying a locking method, one or more indexes, a query-processing operation such as a table scan or index seek, or other options. Table hints are specified in the FROM clause of the DML statement and affect only the table or view referenced in that clause.

This applies to the following commands:

DELETE -

INSERT -

SELECT -

UPDATE -

MERGE -

Reference:

<https://docs.microsoft.com/en-us/sql/t-sql/queries/hints-transact-sql-table>

*Community vote distribution*



🗳️ **NIKnithiy** Highly Voted 1 year, 11 months ago

This is quite tricky WHERE is a clause but SELECT is the statement so answer A is correct  
upvoted 39 times

🗳️ **jordymsft** 1 year, 8 months ago

That is correct.

<https://learn.microsoft.com/en-us/sql/t-sql/statements/statements?view=sql-server-ver16>

<https://data-flair.training/blogs/clause-in-sql/#:~:text=Clases%20are%20in%2Dbuilt%20functions,data%20required%20by%20the%20user.>  
upvoted 4 times

🗳️ **stella\_mah** Highly Voted 1 year, 5 months ago

**Selected Answer: B**

Not a trick question, just how much you understand database.

Select \* From Countries where CountryName = "USA"

This is a full table scan because CountryName is not a key, the engine need to compare every row.

Select \* From Countries where CountryID = 33

This is not a full table scan, it become a index scan because CountryID is a key so the engine know it can skip to that ID

Select \* From Countries

This is a full table scan but with no criteria, so it extract everything out

upvoted 16 times

🗳️ **twfiz** Most Recent 6 months, 2 weeks ago

the correct answer is select. because it is a statement. where is a clause and not a statement

upvoted 1 times

AGTraining 6 months, 2 weeks ago

option A

upvoted 1 times

gabo 7 months ago

Selected Answer: A

In T-SQL, you can use a query hint to instruct the database management system (DBMS) to use a specific index instead of performing a full table scan. The hint you would use is called WITH(INDEX(index\_name)). Here's how you would use it:

```
SELECT column1, column2, ...  
FROM table_name WITH(INDEX(index_name))  
WHERE condition;
```

upvoted 1 times

Yomzie 7 months ago

Answer is option B::

WHERE clause tells a Select to 'streamline' the query execution, on the search criteria (the PREDICATE) specified in the "where" clause.

To use an index in a query, if the index is defined on the column specified in the WHERE clause, the query makes use of the specified index::

```
CREATE TABLE person(  
id INT NOT NULL,  
first_name VARCHAR(15) NOT NULL,  
last_name VARCHAR(15) NOT NULL,  
age INT NOT NULL,  
PRIMARY KEY(id)  
);  
/* Index on First Name */  
CREATE INDEX person_fname ON person (first_name);
```

/\* This next query shows whether the query uses the intended index \*/

```
explain SELECT * FROM person WHERE first_name = "Donald" \G
```

upvoted 1 times

kingAzure 7 months, 2 weeks ago

Selected Answer: A

I agree with NIKnithiy on this one. SELECT is a statement, WHERE is a clause. The link in the answer is also nice as it lists the statements in question. <https://learn.microsoft.com/en-us/sql/t-sql/queries/hints-transact-sql-table?view=sql-server-ver16>

upvoted 1 times

ikramus 7 months, 2 weeks ago

from chatgpt :

To instruct a database management system to use an index instead of performing a full table scan, you typically use a hint in the SELECT statement. Therefore, the correct choice from the provided options is:

A. SELECT

You would use the WITH(INDEX(index\_name)) hint within the SELECT statement to specify the index to be used.

upvoted 1 times

ff09375 8 months, 1 week ago

FWIW, GPT says "B"

upvoted 1 times

Houassa 8 months, 3 weeks ago

Selected Answer: A

The key word here is STATEMENT. WHERE is a CLAUSE.

upvoted 1 times

ravinperera 9 months, 3 weeks ago

A. SELECT

In the SELECT statement, you can specify which index to use for the query. This is typically done using query hints, like WITH

(INDEX(index\_name)), to suggest the use of a particular index. This approach is used when you have a specific understanding of the database's structure and the indexes available, and you want to guide the SQL Server query optimizer to use a specific index. The WHERE and JOIN clauses are used for filtering and combining data from different tables, respectively, and do not directly instruct the use of an index.

upvoted 1 times

🗉 👤 **Law3** 10 months, 1 week ago

SELECT uses the index, correct answer  
SELECT is a statement

WHERE not a statement, it's a clause to filter by column value

```
SELECT * FROM table1 USE INDEX (col1_index,col2_index)
WHERE col1=1 AND col2=2 AND col3=3;
```

upvoted 1 times

🗉 👤 **jaszi** 12 months ago

Which T-SQL statement should be used to instruct a database management system to use an index instead of performing a full table scan?

- A. SELECT
- B. WHERE
- C. JOIN
- ChatGPT
- B. WHERE

The T-SQL statement used to instruct a database management system to use an index instead of performing a full table scan is the WHERE clause in a query. By specifying appropriate conditions in the WHERE clause, you can guide the query optimizer to use available indexes to efficiently retrieve the desired data.

upvoted 2 times

🗉 👤 **Kats** 12 months ago

**Selected Answer: A**

Dropped the exact question in chatgpt, this was the answer...

In T-SQL, you can use the SELECT statement with a query hint to instruct the database management system (DBMS) to use a specific index instead of performing a full table scan. The query hint that accomplishes this is called INDEX or INDEX=<index\_name> hint.

upvoted 1 times

🗉 👤 **AZFabio** 1 year ago

**Selected Answer: A**

"Where" is not a statement, but a clause. Although the filter criteria is defined on "where", it is the "select" statement directs index usage.

upvoted 2 times

🗉 👤 **testdumps2017** 1 year, 1 month ago

it is FROM, but it is not present in the list. so none is correct. I am pretty sure you will get more than 3 options on the exam, FROM will be one of them.

Example: SELECT \* FROM mytable WITH (INDEX (ix\_date)) WHERE field1 > 0

upvoted 1 times

🗉 👤 **HRKK** 1 year, 2 months ago

CORRECT A

upvoted 1 times

HOTSPOT -

Select the answer that correctly completes the sentence.

Hot Area:

### Answer Area

	▼	is a platform as a service (PaaS) relational data service that uses massively parallel processing (MPP).
Azure SQL Database		
Azure Synapse Analytics		
Azure Database for MariaDB		
SQL Server on Azure virtual machines		

### Answer Area

**Suggested Answer:**

	▼	is a platform as a service (PaaS) relational data service that uses massively parallel processing (MPP).
Azure SQL Database		
Azure Synapse Analytics		
Azure Database for MariaDB		
SQL Server on Azure virtual machines		

Azure Synapse Analytics is an platform as a service (PAAS) that combines data integration, warehousing, and analytics into one solution. Azure Synapse Analytics offers cloud-based, relational data warehousing services, massively parallel processing (MPP) scale-out technology, and enough computational power to efficiently manage petabytes and petabytes of data.

Incorrect:

\* Azure Database for MariaDB is a relational database service based on the open-source MariaDB Server engine. It's a fully managed database as a service offering that can handle mission-critical workloads with predictable performance and dynamic scalability.

\* 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. Azure SQL Database is always running on the latest stable version of the SQL Server database engine and patched OS with 99.99% availability. PaaS capabilities built into Azure SQL Database enable you to focus on the domain-specific database administration and optimization activities that are critical for your business.

Reference:

<https://cswsolutions.com/blog/posts/2021/august/what-is-azure-synapse-analytics/> <https://www.integrate.io/blog/what-is-azure-synapse-analytics/>

 **fusioner** Highly Voted 2 years ago

Correct Answer

upvoted 15 times

 **JA2018** Highly Voted 1 year, 11 months ago

Answer: B

<https://learn.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/massively-parallel-processing-mpp-architecture>

upvoted 8 times

 **b94da7f** Most Recent 4 months, 1 week ago

Azure Synapse analytics

upvoted 1 times

 **AGTraining** 6 months, 2 weeks ago

option B

upvoted 1 times

 **GaneshNats** 1 year, 4 months ago

Azure Synapse analytics

upvoted 4 times

 **chiiiweiii** 1 year, 6 months ago

Correct Answer

upvoted 2 times

  **blackskies** 1 year, 11 months ago

was on exam today 09/30/22

upvoted 4 times

HOTSPOT -

Select the answer that correctly completes the sentence.

Hot Area:

### 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		

### Answer Area

**Suggested Answer:**

	▼	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 clustered index defines the order in which data is physically stored in a table. Table data can be sorted in only way, therefore, there can be only one clustered index per table. In SQL Server, the primary key constraint automatically creates a clustered index on that particular column.

Reference:

<https://www.sqlshack.com/what-is-the-difference-between-clustered-and-non-clustered-indexes-in-sql-server>

 **fusioner** Highly Voted 2 years ago

Correct Answer

upvoted 10 times

 **raydel92** Highly Voted 1 year, 11 months ago

Answer is correct.

Clustered indexes sort and store the data rows in the table or view based on their key values. These are the columns included in the index definition.

Ref: <https://learn.microsoft.com/en-us/sql/relational-databases/indexes/clustered-and-nonclustered-indexes-described?view=sql-server-ver16>

upvoted 7 times

 **b94da7f** Most Recent 4 months, 1 week ago

Clustered indexes

upvoted 1 times

 **AGTraining** 6 months, 2 weeks ago

option B

upvoted 1 times

 **Edgy\_San** 7 months, 1 week ago

What helps me remember things related to synapse, spark (firing of neurons) is remembering the brain. MPP- distributed computer nodes working parallel just like a network of neurons in the brain as well.

upvoted 2 times

 **testdumps2017** 1 year, 1 month ago

Clustered index will NOT sort the data physically. I know they want you to select that answer, but if you ever go for an interview, don't even think of saying "data is stored physically by the clustering key".

So, clustered index is correct for the exam, incorrect in real-life.

<https://sqlwithmanoj.com/2013/06/02/clustered-index-do-not-guarantee-physically-ordering-or-sorting-of-rows/>

upvoted 1 times

 **chiiiweiii** 1 year, 6 months ago

Correct Answer

upvoted 2 times

🗨️ 👤 **Pinha** 1 year, 7 months ago

Correct Answer.

if you have a table with a clustered index on a column called "id", all the rows with the lowest "id" value will be stored together physically on disk, followed by all the rows with the next lowest "id" value, and so on. This physical sorting allows the database to quickly locate and retrieve specific rows of data based on their indexed column values, improving query performance.

upvoted 2 times

🗨️ 👤 **Nurialzard** 1 year, 8 months ago

same question as number 40

upvoted 3 times

HOTSPOT -

Select the answer that correctly completes the sentence.

Hot Area:

### Answer Area

In a relational database, the  uniquely identifies each row in a table

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

### Answer Area

Suggested Answer:

In a relational database, the  uniquely identifies each row in a table

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

A primary key is a special relational database table column (or combination of columns) designated to uniquely identify each table record (row).

Reference:

<https://www.techopedia.com/definition/5547/primary-key>

  **fusioner** Highly Voted 2 years ago

Primary key

upvoted 16 times

  **b94da7f** Most Recent 4 months, 1 week ago

Primary key

upvoted 1 times

  **AGTraining** 6 months, 2 weeks ago

option B

upvoted 1 times

  **chiiiweiii** 1 year, 6 months ago

Correct Answer

upvoted 4 times

  **Forecaster** 1 year, 9 months ago

Answer is correct! Its B, Primary key

upvoted 4 times

HOTSPOT -

Select the answer that correctly completes the sentence.

Hot Area:

### Answer Area

	▼	presents content defined by a query.
A heap		
A stored procedure		
A view		
An index		

### Answer Area

Suggested Answer:

	▼	presents content defined by a query.
A heap		
A stored procedure		
A view		
An index		

A database view is a subset of a database and is based on a query that runs on one or more database tables.

Reference:

[https://www.ibm.com/docs/SSWT9A\\_7.6.0/com.ibm.mbs.doc/configur/c\\_views.html](https://www.ibm.com/docs/SSWT9A_7.6.0/com.ibm.mbs.doc/configur/c_views.html)

- 🗄️ **fusioner** Highly Voted 2 years ago  
 Repeated Question  
 upvoted 9 times
- 🗄️ **PatYeo** Highly Voted 1 year, 11 months ago  
 View is correct.  
 upvoted 7 times
- 🗄️ **b94da7f** Most Recent 4 months, 1 week ago  
 A VIEW  
 upvoted 1 times
- 🗄️ **sjokkorein** 5 months, 4 weeks ago  
 quarry = view  
 upvoted 2 times
- 🗄️ **AGTraining** 6 months, 2 weeks ago  
 option C  
 upvoted 1 times
- 🗄️ **akhil5432** 1 year, 1 month ago  
 A VIEW  
 upvoted 1 times
- 🗄️ **mariem50** 1 year, 2 months ago  
 Correct Answer  
 upvoted 1 times
- 🗄️ **chiiiweiii** 1 year, 6 months ago  
 Correct Answer  
 upvoted 3 times
- 🗄️ **Nurialzard** 1 year, 8 months ago  
 question 43  
 upvoted 1 times

Which Azure service provides the highest compatibility for databases migrated from Microsoft SQL Server 2019 Enterprise edition?

- A. Azure SQL Database
- B. Azure Database for MySQL
- C. Azure SQL Managed Instance
- D. an Azure SQL Database elastic pool

**Suggested Answer:** C

SQL Managed Instance has near 100% compatibility with the latest SQL Server (Enterprise Edition) database engine, providing a native virtual network (VNet) implementation that addresses common security concerns, and a business model favorable for existing SQL Server customers.

Note: Azure SQL Managed Instance is the intelligent, scalable cloud database service that combines the broadest SQL Server database engine compatibility with all the benefits of a fully managed and evergreen platform as a service.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/managed-instance/sql-managed-instance-paas-overview?view=azuresql>

Community vote distribution

C (100%)

 **john1979** Highly Voted 1 year, 11 months ago

i am with C

upvoted 11 times

 **mamahani** Highly Voted 1 year, 5 months ago

I took the exam 12.04.2023; this question appeared, correct answer:

Azure SQL Managed Instance

upvoted 10 times

 **okayhey** Most Recent 6 months, 1 week ago

Had this one on the exam, 8-3-2024!

upvoted 1 times

 **AGTraining** 6 months, 2 weeks ago

option C

upvoted 1 times

 **AZFabio** 1 year ago

Selected Answer: C

SQL Managed Instances

upvoted 1 times

 **choby** 1 year, 4 months ago

For me: C.Azure SQL Managed Instance

upvoted 1 times

 **chiiweiii** 1 year, 6 months ago

Selected Answer: C

Correct Answer

upvoted 3 times

 **BeliteLiu** 1 year, 7 months ago

Selected Answer: C

ANSER IS C

upvoted 2 times

 **[Removed]** 1 year, 9 months ago

Selected Answer: C

C. Azure SQL MI is correct.

upvoted 4 times

🗨️ 👤 **cliffzx** 1 year, 9 months ago

Answer is C

upvoted 2 times

🗨️ 👤 **mantran** 1 year, 10 months ago

C is correct

upvoted 4 times

🗨️ 👤 **JA2018** 1 year, 11 months ago

**Selected Answer: C**

Answer: C

<https://learn.microsoft.com/en-us/azure/azure-sql/managed-instance/sql-managed-instance-paas-overview?view=azuresql>

upvoted 4 times

🗨️ 👤 **fusioner** 2 years ago

Correct Answer C

upvoted 4 times

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-production account type
- D. API

**Suggested Answer: D**

You must create a separate account for each API.

Reference:

<https://docs.microsoft.com/en-us/azure/cosmos-db/how-to-manage-database-account>

Community vote distribution

D (85%)

Other

 **confetti** Highly Voted 2 years, 5 months ago

D is the answer. Verified on the portal

upvoted 16 times

 **Bambi0074** Highly Voted 1 year, 11 months ago

**Selected Answer: D**

You must choose the API so you can't switch this within the creation process.

Got this description from the creation page:

Create Azure Cosmos DB Account - Core (SQL)

Global Distribution

"Configure global distribution and regional settings for your account. You can also change these settings after the account is created."

upvoted 10 times

 **srgopalkr** 1 year, 1 month ago

I tried creating an Azure Cosmos DB for NoSQL and didn't find the option to choose the API. However, I'm able to see the 'API' parameter being auto set to 'Azure Cosmos DB for SQL' in the 'Review+Create' page. Perhaps this was a configurable attribute in the past, but is no more.

upvoted 2 times

 **AGTraining** Most Recent 6 months, 2 weeks ago

option D

upvoted 1 times

 **rmeng** 1 year, 1 month ago

**Selected Answer: C**

C. production or non-production account type

The setting that can only be configured during the creation of an Azure Cosmos DB account is the account type, specifically choosing between a production or non-production account.

upvoted 1 times

 **Sr18** 4 months, 2 weeks ago

Are you serious!!! Answer is D

its just sets a tag value while creating cosmos, that later can be changed to non production anytime. look for CosmosAccountType

upvoted 1 times

 **srgopalkr** 1 year, 1 month ago

I tried creating an Azure Cosmos DB for NoSQL and don't find this option (production or non-production account type) listed. As they keep updating these configs on a regular basis, perhaps this option existed in the past but NOT available anymore.

upvoted 2 times

 **gabo** 7 months ago

Not sure when this option existed but not anymore.

upvoted 1 times

🗨️ **chiiweiii** 1 year, 6 months ago

Selected Answer: D

Correct Answer  
upvoted 3 times

🗨️ **Joanna0** 1 year, 7 months ago

Selected Answer: D

Setting: API

Description: In this, select Core (SQL) for creating a document database and query by using SQL syntax. However, the API determines the type of account to create. Further, Azure Cosmos DB provides five APIs that include Core (SQL) and MongoDB for document data, Gremlin for graph data, Azure Table, and Cassandra.

upvoted 2 times

🗨️ **JanJandnb** 1 year, 11 months ago

Geo-Redundancy Disable Enable or disable global distribution on your account by pairing your region with a pair region. You can add more regions to your account later.

upvoted 1 times

🗨️ **srgopalkr** 1 year, 1 month ago

I tried creating an 'Azure Cosmos DB for NoSQL' and in the 'Global Distribution' page of the creation process, I do see the settings for 'Geo-Redundancy' & 'Multi-Region Writes', but there is also a disclaimer note that "You can also change these settings after the account is created.". However, the question is about a setting that can be configured only during the account creation process and NOT thereafter.

upvoted 2 times

🗨️ **himanshithi55** 1 year, 12 months ago

Currently, you must create a separate account for each API

upvoted 1 times

🗨️ **rohjay** 2 years ago

The answer should be "A" as per the following explanation in point no 5

<https://docs.microsoft.com/en-us/azure/cosmos-db/how-to-manage-database-account>

5. In the Global Distribution tab, configure the following details. You can leave the default values for this quickstart:

option available :Geo-Redundancy and Multi-region Writes

upvoted 1 times

🗨️ **pimni** 2 years, 4 months ago

D is the answer

upvoted 6 times

🗨️ **bfan95** 2 years, 4 months ago

Selected Answer: D

Answer: D

upvoted 3 times

🗨️ **\_xXx\_** 2 years, 5 months ago

Selected Answer: A

API can only be set on account level when creating an account

upvoted 3 times

🗨️ **gaem** 2 years, 6 months ago

Selected Answer: D

D is the answer

upvoted 4 times

🗨️ **linyuyouxiang** 2 years, 9 months ago

API must be specified

upvoted 2 times

🗨️ **Ghibli** 2 years, 9 months ago

Yes one can set both API as well as production or non-reduction account type while creating account. I am confused with this question, can anyone explain?

upvoted 2 times

🗨️ **BabieTee** 2 years, 3 months ago

The following options are not available if you select Serverless as the Capacity mode:

Apply Free Tier Discount

Geo-redundancy

Multi-region Writes

upvoted 1 times

  **aryoyo** 2 years, 10 months ago

answer is C: production or non-reduction account type

upvoted 2 times

  **aryoyo** 2 years, 10 months ago

Correction I agree with API (the focus in this question is 'during creation of account')

upvoted 3 times

  **Marskmen** 2 years, 10 months ago

Answer is correct

upvoted 3 times

DRAG DROP -

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.

Select and Place:

### Data Store Types

- Graph
- Object
- Key/value

### Answer Area

A list of application user names and, for each, the default language

A collection of medical images and their associated metadata.

Employee data that shows the relationships between employees

### Suggested Answer:

#### Data Store Types

- Graph
- Object
- Key/value

#### Answer Area

- Key/value
- Object
- Graph

A list of application user names and, for each, the default language

A collection of medical images and their associated metadata.

Employee data that shows the relationships between employees

Box 1: Key/value -

A key/value store associates each data value with a unique key.

Box 2: Object -

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).

Box 3: Graph -

A graph database stores two types of information, nodes and edges. Edges specify relationships between nodes. Nodes and edges can have properties that provide information about that node or edge, similar to columns in a table. Edges can also have a direction indicating the nature of the relationship.

Reference:

<https://docs.microsoft.com/en-us/azure/architecture/guide/technology-choices/data-store-overview>

 **blackskies** Highly Voted 1 year, 11 months ago

Correct

upvoted 22 times

 **AGTraining** Most Recent 6 months, 2 weeks ago

Key/value

object

graph

upvoted 2 times

 **[Removed]** 1 year, 4 months ago

Key value, object and graph indeed!

upvoted 2 times

 **Ejake** 1 year, 5 months ago

Key/value

object

graph

upvoted 4 times

  **chiiweiii** 1 year, 6 months ago

Correct Answer

upvoted 3 times

  **Yamarh** 1 year, 9 months ago

Correct

upvoted 3 times

  **Sheoran21** 1 year, 9 months ago

CORRECT -

upvoted 4 times

You have an Azure Cosmos DB account that uses the Core (SQL) API.

Which two settings can you configure at the container level? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. the throughput
- B. the read region
- C. the partition key
- D. the API

**Suggested Answer:** AC

Reference:

<https://www.sqlshack.com/start-your-journey-with-azure-cosmos-db/>

Community vote distribution

AC (100%)

- 🗨️ **PatYeo** Highly Voted 1 year, 11 months ago  
A and C are correct : <https://www.sqlshack.com/wp-content/uploads/2020/04/create-database-container.png>  
upvoted 17 times
- 🗨️ **srgopalkr** Highly Voted 1 year, 1 month ago  
Please note that the Azure Cosmos DB set-up screens have changed, making this question non-relevant. The link referenced in the answer was a screenshot taken in 2020.  
upvoted 9 times
- 🗨️ **AGTraining** Most Recent 6 months, 2 weeks ago  
option A C  
upvoted 2 times
- 🗨️ **chiiiweiii** 1 year, 6 months ago  
Selected Answer: AC  
Correct Answers  
upvoted 3 times
- 🗨️ **Joanna0** 1 year, 7 months ago  
Selected Answer: AC  
<https://learn.microsoft.com/en-us/azure/cosmos-db/account-databases-containers-items#azure-cosmos-db-containers>  
upvoted 6 times
- 🗨️ **captain\_60** 1 year, 7 months ago  
A and c is correct  
upvoted 2 times
- 🗨️ **alex\_0207** 1 year, 8 months ago  
AC is correct  
upvoted 2 times
- 🗨️ **PatYeo** 1 year, 11 months ago  
A and C is correct : <https://docs.microsoft.com/en-us/azure/cosmos-db/how-to-manage-database-account>  
upvoted 7 times
- 🗨️ **fusioner** 2 years ago  
Selected Answer: AC  
AC is Correct  
upvoted 4 times
- 🗨️ **Mern** 2 years, 3 months ago  
Correct  
upvoted 4 times

  **bfan95** 2 years, 4 months ago

AC is correct

upvoted 4 times

  **svetlanam** 2 years, 4 months ago

Correct.

upvoted 5 times

Your company is designing a data store that will contain student data. The data has the following format.

StudentNumber	StudentInformation
7634634	First name: Ben Last: Smith Preferred Name: Benjamin
7634634	First Name: Dominik Last Name: Paiha Email Address: dpaiha@contoso.com MCP ID: 931817
7634636	First Name: Reshma Last Name: Patel Phone number: 514-555-1101
7634637	First Name: Yun-Feng Last Name: Peng

Which type of data store should you use?

- A. graph
- B. key/value
- C. object
- D. columnar

**Suggested Answer:** D

Community vote distribution



**blablax** Highly Voted 4 years ago

<https://docs.microsoft.com/en-us/azure/architecture/guide/technology-choices/data-store-overview#column-family-databases>  
upvoted 53 times

**sourabh7257** 4 years ago

are answers correct?  
upvoted 4 times

**Dhiva\_** 4 years ago

Yes it is  
upvoted 17 times

**zxmaz01** Highly Voted 2 years, 3 months ago

Am I the only one, who is confused by the first three students having the same number? :-D  
upvoted 19 times

**AGTraining** Most Recent 6 months, 2 weeks ago

option D  
upvoted 1 times

**AZFabio** 1 year ago

Selected Answer: D

Columnar as you can handle Student Info as a column family.  
upvoted 2 times

**rmeng** 1 year ago

Selected Answer: C

C. object

The data seems to fit the structure of an object or document-oriented database, where each student's information can be stored as a document

with various attributes and their corresponding values. This type of data store allows for flexible schema and can accommodate varying attributes for different documents.

upvoted 2 times

  **gabo** 7 months ago

I thought the same - each student can be stored as a JSON Document in an Azure Cosmos DB

upvoted 1 times

  **srgopalkr** 1 year, 1 month ago

Answer is D.

Here's a useful link to help visualize columnar database.

upvoted 1 times

  **srgopalkr** 1 year, 1 month ago

Missed to include the link - <https://www.kdnuggets.com/2021/02/understanding-nosql-database-types-column-oriented-databases.html>

upvoted 1 times

  **rmeng** 1 year, 1 month ago

**Selected Answer: C**

C. object

Based on the given format of student data, an object data store would be the most suitable choice.

In an object data store, data is stored as objects, where each object is a self-contained unit consisting of both data and associated metadata.

The data within an object can be structured or unstructured, and it can contain various attributes and properties.

In the given scenario, each student's data is presented as an object with attributes such as "First Name," "Last Name," "Email Address," "MCP ID," and "Phone Number." The data is self-contained within each object, making it an ideal fit for an object data store.

The object data store is designed for scenarios where you need to store complex, heterogeneous data with varying attributes and metadata.

upvoted 1 times

  **pirpa** 1 year, 6 months ago

**Selected Answer: A**

Correct Answer

upvoted 3 times

  **pirpa** 1 year, 6 months ago

I've selected the wrong option. The correct answer is D. Moderator could you change my selected answer from A to B. Thanks

upvoted 1 times

  **SrishtiShetty** 6 months, 2 weeks ago

Are you confused? Your answer is D or B? You have mentioned correct answer is D but then u have also mentioned the moderator to change A to B

upvoted 1 times

  **chiiweiii** 1 year, 6 months ago

Correct Answer

upvoted 2 times

  **Joanna0** 1 year, 7 months ago

**Selected Answer: D**

<https://learn.microsoft.com/en-us/azure/architecture/guide/technology-choices/data-store-overview#column-family-databases>

upvoted 1 times

  **jgsa11** 1 year, 7 months ago

key value is a possible answer since the second column is aggregating like in a JSON with varying subschema.

upvoted 5 times

  **zerone72** 1 year ago

I guess it's not a key-value because the first 2 students have same number

upvoted 1 times

  **JustinYoo** 1 year, 8 months ago

The 2 given student numbers in the table are the same but different name? the question is wrong.

upvoted 3 times

  **Dusica** 1 year, 9 months ago

**Selected Answer: D**

A, B, C are out of question; so, it must be D. However Columnar database is not same as Column-family data store, which this exhibit IS. So, who named it Columnar: Microsoft on the exam or those who retrieved the question from their memory?

upvoted 5 times

  **Dusica** 1 year, 10 months ago

A, B, C are out of question; so, it must be D. However, Columnar database is not same as Column-family data store, which this exhibit IS. So, who named it Columnar: Microsoft on the exam or those who retrieved the question from their memory?

upvoted 1 times

  **mayconze** 1 year, 11 months ago

**Selected Answer: D**

The answer is correct.

upvoted 2 times

  **fusioner** 2 years ago

**Selected Answer: D**

Columnar

upvoted 3 times

  **Ayor** 2 years, 3 months ago

Answer is D

upvoted 2 times

Which storage solution supports role-based access control (RBAC) at the file and folder level?

- A. Azure Disk Storage
- B. Azure Data Lake Storage
- C. Azure Blob storage
- D. Azure Queue storage

**Suggested Answer:** B

Reference:

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

Community vote distribution

B (100%)

 **Dhiva\_** Highly Voted 4 years ago

The given answer is correct  
upvoted 37 times

 **GetulioJr** Highly Voted 3 years, 3 months ago

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). This article summarizes the basics of the access control model for Data Lake Storage Gen2.

ACLs are applied on the file and folder level. The key thing to remember is that you are always going to need RBAC Control Plane permissions in combination with ACLs. Best practice is to assign your security principals RBAC Reader role on the Storage Account/Container level and continue with more restrictive ACLs on the file and folder level.

There are two types of ACLs:

Access ACLs: They control access to an object. An object can be a file or a folder.

Default ACLs: These are ACLs assigned on the folder level only which get inherited as Access ACLs by the child file or folder.

upvoted 30 times

 **AGTraining** Most Recent 6 months, 2 weeks ago

option B  
upvoted 1 times

 **AZFabio** 1 year ago

Selected Answer: B

Data Lake Storage is correct  
upvoted 1 times

 **akhil5432** 1 year, 1 month ago

Selected Answer: B

AZ data alke  
upvoted 1 times

 **chiiiweiii** 1 year, 6 months ago

Selected Answer: B

Correct Answer  
upvoted 2 times

 **Joanna0** 1 year, 7 months ago

Selected Answer: B

<https://learn.microsoft.com/en-us/azure/storage/blobs/data-lake-storage-access-control-model#role-based-access-control-azure-rbac>  
upvoted 2 times

 **Drchattss** 2 years, 1 month ago

01/08.....EXAM

upvoted 4 times

🗨️ **somsunu** 2 years, 3 months ago

correct

upvoted 3 times

🗨️ **print>HelloWorld** 2 years, 4 months ago

CORRECT

upvoted 2 times

🗨️ **pimni** 2 years, 4 months ago

correct

upvoted 2 times

🗨️ **Xuitex** 2 years, 4 months ago

Correct

upvoted 2 times

🗨️ **roshanazMaster** 2 years, 4 months ago

**Selected Answer: B**

correct

upvoted 2 times

🗨️ **Teamdev** 2 years, 5 months ago

for sure, it's correct

upvoted 2 times

🗨️ **SiDoCiOuS** 2 years, 5 months ago

**Selected Answer: B**

Appeared on the exam on 03/22/2022

Answer is correct!

upvoted 4 times

🗨️ **Nikhil1702** 2 years, 7 months ago

B IS CORRECT

upvoted 2 times

🗨️ **sounakroy** 2 years, 7 months ago

**Selected Answer: B**

Correct

upvoted 2 times

You need to store data in Azure Blob storage for seven years to meet your company's compliance requirements. The retrieval time of the data is unimportant. The solution must minimize storage costs.

Which storage tier should you use?

- A. Archive
- B. Hot
- C. Cool

**Suggested Answer: A**

Reference:

[https://cloud.netapp.com/blog/azure-blob-storage-pricing-the-complete-guide-azure-cvo-blg#H1\\_4](https://cloud.netapp.com/blog/azure-blob-storage-pricing-the-complete-guide-azure-cvo-blg#H1_4)

Community vote distribution

A (100%)

 **Dhiva\_** Highly Voted 4 years ago

The answer is correct

<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-storage-tiers?tabs=azure-portal>

upvoted 34 times

 **User6** Highly Voted 3 years, 9 months ago

Answer is correct Hot - Optimized for storing data that is accessed frequently.

Cool - Optimized for storing data that is infrequently accessed and stored for at least 30 days.

Archive - Optimized for storing data that is rarely accessed and stored for at least 180 days with flexible latency requirements (on the order of hours).

upvoted 21 times

 **Luis\_polanco** 2 years, 6 months ago

that is not correct, it says that The retrieval time of the data is unimportant. So A is the correct answer, it is looking for saving costs

upvoted 3 times

 **AGTraining** Most Recent 6 months, 2 weeks ago

option A

upvoted 1 times

 **Rudramina** 8 months, 1 week ago

Archive Tier

upvoted 1 times

 **chiiweiii** 1 year, 6 months ago

Selected Answer: A

Correct Answer

upvoted 4 times

 **JA2018** 1 year, 11 months ago

Selected Answer: A

the only logical choice

upvoted 2 times

 **gabrisiq** 2 years, 1 month ago

definitely Archive

upvoted 3 times

 **RichMelbourne** 2 years, 5 months ago

Shown on test 15/04/2022

upvoted 2 times

 **Teamdev** 2 years, 5 months ago

A is the correct answer

upvoted 1 times

🗨️ 👤 **tittu1947** 2 years, 5 months ago

The answer is correct

upvoted 1 times

🗨️ 👤 **D\_Mo\_ZY** 2 years, 6 months ago

Correct

Was in the exam (18/03/2022)

upvoted 3 times

🗨️ 👤 **WigglyPork** 2 years, 7 months ago

Retrieval Rate:

Hot - Frequent

Cold - every 30 days

Archive - every 180 days

upvoted 8 times

🗨️ 👤 **Koph** 2 years, 9 months ago

Archive

upvoted 1 times

🗨️ 👤 **jerryi** 3 years ago

A is the correct Answer

Archive - Optimized for storing data that is rarely accessed and stored for at least 180 days with flexible latency requirements, on the order of hours.

upvoted 3 times

🗨️ 👤 **mpawan** 3 years, 8 months ago

Correct,. Archive it to save money

upvoted 1 times

🗨️ 👤 **EBAN** 3 years, 11 months ago

OK, but remember that in Archive Storage, after rehydrating, you pay FOR ACCESS, so depending on how much access you intend to do, it can be very expensive. In my Opinion and taking a look at other trainings, it should be Cool, that is less expensive than Hot, more expensive than Archive but you don't pay for access. This answer can be right only if you think that the number of access is not declared.

upvoted 6 times

🗨️ 👤 **klever** 3 years, 8 months ago

I think the key piece of information is: The retrieval time of the data is unimportant. So archive is correct answer

upvoted 5 times

🗨️ 👤 **msbj** 2 years, 9 months ago

agreed

upvoted 2 times

🗨️ 👤 **Razh** 3 years, 9 months ago

I would have gone for 'cool'. The question itself is quite dull as in not much detail...

upvoted 2 times

Which type of non-relational data store supports a flexible schema, stores data as JSON files, and stores the all the data for an entity in the same document?

- A. document
- B. columnar
- C. graph
- D. time series

**Suggested Answer: A**

Community vote distribution

A (100%)

 **RicardoLessa** Highly Voted 3 years, 12 months ago

The given answer is correct!!  
upvoted 45 times

 **vhx** Highly Voted 3 years, 11 months ago

Yes it is correct  
upvoted 12 times

 **AGTraining** Most Recent 6 months, 2 weeks ago

option A  
upvoted 1 times

 **AZFabio** 1 year ago

Selected Answer: A  
document  
upvoted 1 times

 **mamahani** 1 year, 5 months ago

I took the exam 12.04.2023; this question appeared, correct answer:  
document  
upvoted 8 times

 **chiiiweiii** 1 year, 6 months ago

Selected Answer: A  
Correct Answer  
upvoted 2 times

 **facilitator** 1 year, 6 months ago

A is correct  
upvoted 2 times

 **KMLearn3** 1 year, 9 months ago

Very similar to Qst#110: the example there was not relational, no fix schema. etc.  
Despite of that "Document" was false there.  
upvoted 2 times

 **TomGray1989** 1 year, 10 months ago

This question blows my mind...but the answer is correct. Document.  
upvoted 1 times

 **Manoj6** 2 years, 3 months ago

Document is correct  
<https://docs.microsoft.com/en-us/azure/architecture/guide/technology-choices/data-store-overview#column-family-databases>  
upvoted 2 times

 **Teamdev** 2 years, 5 months ago

Yes correct

upvoted 1 times

🗨️ 👤 **Mopo** 2 years, 5 months ago

**Selected Answer: A**

Correct

upvoted 2 times

🗨️ 👤 **olitkx** 2 years, 8 months ago

**Selected Answer: A**

correct

upvoted 5 times

🗨️ 👤 **Koph** 2 years, 9 months ago

Document is correct

upvoted 2 times

🗨️ 👤 **[Removed]** 3 years, 3 months ago

answer is correct. the key is json = document

upvoted 8 times

🗨️ 👤 **mnpawan** 3 years, 8 months ago

A. Document DB

upvoted 2 times

🗨️ 👤 **IAMKPR** 3 years, 10 months ago

[https://en.wikipedia.org/wiki/Document-oriented\\_database](https://en.wikipedia.org/wiki/Document-oriented_database)

upvoted 4 times

DRAG DROP -

Match the Azure Cosmos DB APIs to the appropriate data structures.

To answer, drag the appropriate API from the column on the left to its data structure on the right. Each API may be used once, more than once, or not at all.

NOTE: Each correct match is worth one point.

Select and Place:

**APIs**

- Cassandra API
- Gremlin API
- MongoDB API
- Table API

**Answer Area**

- 
- 
- Key/value data

**APIs**

- Cassandra API
- Gremlin API
- MongoDB API
- Table API

**Answer Area**

- Gremlin API
- MongoDB API
- Table API

Reference:  
<https://docs.microsoft.com/en-us/azure/cosmos-db/faq>

- ☰ **Average\_Joe** Highly Voted 3 years, 3 months ago  
 GREmlin  
 GRaph  
 upvoted 44 times
- ☰ **\_Noe\_** 2 years, 8 months ago  
 Yes, and also  
 C assandra  
 C olumnar  
 upvoted 28 times
- ☰ **BieLey** 2 years, 3 months ago  
 Table  
 Keyboard  
 upvoted 14 times
- ☰ **harshgulhane** Highly Voted 3 years, 11 months ago  
 Yes, I got this question in exam  
 upvoted 41 times
- ☰ **AGTraining** Most Recent 6 months, 2 weeks ago  
 Gremlin API  
 MongoDB API  
 Table API  
 upvoted 1 times
- ☰ **chiiiweiii** 1 year, 6 months ago  
 Correct Answers  
 upvoted 4 times
- ☰ **facilitator** 1 year, 6 months ago

The Gremlin API is used for graph databases. The MongoDB API stores data in the BSON format. The Table API is used to retrieve key/value pairs. The Cassandra API is used to retrieve tabular data.

upvoted 7 times

🗨️ 👤 **RichMelbourne** 2 years, 5 months ago

Shown on test 15/04/2022

upvoted 5 times

🗨️ 👤 **Teamdev** 2 years, 5 months ago

giving answer is correct

upvoted 1 times

🗨️ 👤 **SiDoCiOuS** 2 years, 5 months ago

Appeared on the exam on 03/22/2022

Answers are correct! Gremlin, Mongo, Table

upvoted 2 times

🗨️ 👤 **mrjeet** 2 years, 8 months ago

Gremlin API - This API allows users to make graph queries and stores data as edges and vertices

MongoDB API - This API stores data in a document structure, via BSON format

Table API - This API stores data in key/value format

upvoted 9 times

🗨️ 👤 **Jeetesh** 2 years, 5 months ago

Which means that the 2nd answer provided is wrong. Gremlin API used JSON format

upvoted 1 times

🗨️ 👤 **charuu71** 2 years, 9 months ago

Mongo API : This API stores data in a document structure, via BSON format.

Cassandra API : This API stores data in column-oriented schema.

upvoted 2 times

🗨️ 👤 **76Carpediem** 2 years, 9 months ago

It's correct

upvoted 3 times

🗨️ 👤 **nadjar007** 2 years, 9 months ago

It's correct. <https://docs.microsoft.com/en-us/azure/cosmos-db/choose-api>

upvoted 2 times

🗨️ 👤 **Meyti** 2 years, 9 months ago

...MongoDB is another well-known Document database... also ----> JSON files

upvoted 2 times

🗨️ 👤 **theLazyGuy** 2 years, 10 months ago

Gremlin API...allows users to make graph queries

Table API stores data in key/value format

(<https://docs.microsoft.com/en-us/azure/cosmos-db/choose-api>)

upvoted 2 times

🗨️ 👤 **GetulioJr** 3 years, 3 months ago

Column-Family Data Model:

This data model supports Cassandra API. Existing Cassandra implementations can be easily and quickly moved to Cosmos DB and use column family format, which is used in Cassandra. Column Family data model is similar to key-value data model except that the items in the data model adhere to the defined schema.

Key-Value Pair Data Model (Table API):

In this data model, each entity consists of a key and a value pair. The value itself can be a set of key-value pairs. This is very similar to a table in a relational database where each row has the same set of columns. The Key-Value pair solution is supported on standard Azure table API. This Table API is primarily helpful for existing Azure Table Storage customers in migration to Cosmos DB.

REF: <https://www.red-gate.com/simple-talk/cloud/azure/overview-of-azure-cosmos-db/>

upvoted 4 times

🗨️ 👤 **zwy** 3 years, 3 months ago

Column family HBase in HDInsight and Cosmos DB (using Cassandra API)

so the given answer is correct

upvoted 1 times

  **pac1311** 3 years, 5 months ago

Table API is a part of azure storage service, no?

upvoted 1 times

HOTSPOT -

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

Hot Area:

### 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).

Suggested Answer:

### 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 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. With a hierarchical namespace enabled, a storage account becomes capable of providing the scalability and cost-effectiveness of object storage, with file system semantics that are familiar to analytics engines and frameworks.

One advantage of hierarchical namespace is atomic directory manipulation: Object stores approximate a directory hierarchy by adopting a convention of embedding slashes (/) in the object name to denote path segments.

Reference:

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

 **IAMKPR** Highly Voted 3 years, 10 months ago

Answer is correct and able to understand with given reference article.

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

upvoted 32 times

 **Pinha** Highly Voted 1 year, 7 months ago

Security at the folder level: With hierarchical namespace enabled, you can set different security permissions (e.g. read, write, delete) on different folders within the storage account, allowing you to better control access to your data.

Atomic directory manipulation: With hierarchical namespace enabled, you can perform atomic operations on directories, such as renaming or moving a directory and all its contents in one single operation. This ensures that the entire operation is completed or rolled back as a whole, avoiding partial changes to the data.

upvoted 9 times

 **AGTraining** Most Recent 6 months, 2 weeks ago

option A

upvoted 1 times

 **chiiiweiii** 1 year, 6 months ago

Correct Answer

upvoted 5 times

 **Ashishsingh07** 2 years, 4 months ago

FEATURE IS VALID FOR ADLS , WHICH IS ENABLED USING HIERARCHIAL NAMESPACE :)

upvoted 3 times

🗨️ 👤 **RichMelbourne** 2 years, 5 months ago

Shown on test 15/04/2022

upvoted 2 times

🗨️ 👤 **Teamdev** 2 years, 5 months ago

Correct!

upvoted 3 times

🗨️ 👤 **vcehunger** 2 years, 7 months ago

Answer is correct

upvoted 3 times

🗨️ 👤 **Phong0411** 2 years, 9 months ago

Answer is correct

upvoted 2 times

🗨️ 👤 **Meyti** 2 years, 9 months ago

A tricky question:)

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).

To create Azure Storage account for Data Lake Storage you must enable hierarchal namespace.

Then Answer is A.

upvoted 4 times

🗨️ 👤 **PlakaPlaka** 2 years, 10 months ago

Answer is correct

upvoted 1 times

🗨️ 👤 **apgl** 3 years, 8 months ago

Data Lake support POSIX filesystem, it allows to manage roles and rights at folder/file level, at the same manner as in Unix systems.

upvoted 4 times

🗨️ 👤 **harshgulhane** 3 years, 11 months ago

Tricky question, confused me alot.

upvoted 3 times

🗨️ 👤 **Preben** 3 years, 6 months ago

I believe the key in this one was to catch the word "folder" in the text. Only hierarchal namespace-based storages (like Data Lake) uses folders.

upvoted 9 times

HOTSPOT -

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

Hot 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.

**Suggested Answer:**

**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.

Reference:

<https://www.sqlshack.com/graph-database-implementation-with-azure-cosmos-db-using-the-api/>

 **jasu** Highly Voted 4 years ago

Wrong answer is given. It should be Gremlin API. <https://docs.microsoft.com/en-us/azure/cosmos-db/graph-introduction>  
upvoted 203 times

 **vhx** 3 years, 11 months ago

Definitely D is the answer  
upvoted 28 times

 **casualcrackers** 2 years, 5 months ago

i had this question today and D was the answer A was not an option  
upvoted 29 times

 **MinimizeEfforts** 2 years, 7 months ago

Be careful! This highly voted comment is misleading The question is tricky and the correct answer is definately "B" (use SQL-like statement to query a graph db)  
upvoted 13 times

 **EBAN** Highly Voted 3 years, 11 months ago

Attention folks. The question is how you query and not what is needed to connect Graph to Cosmos DB. Once is connected, you use Cosmos DB to perform queries. In my HUMBLE opinion, the given answer is right  
upvoted 51 times

 **Omar\_Alfatih** 3 years, 3 months ago

here is the link for running queries... its Gremlin <https://docs.microsoft.com/en-us/azure/cosmos-db/tutorial-query-graph>  
upvoted 5 times

 **Abhishek\_89** 2 years, 3 months ago

I had the question today in test, the option .. ' as json doc..' was not an option.  
upvoted 8 times

 **lebigmac** 3 years, 7 months ago

You are correct! The link provided, at the very bottom, has a very clear example, that you use a SQL-like statement to query a graph db and returns a JSON doc as results. The answer is correct.  
upvoted 5 times

 **Abhishek17** Most Recent 2 months, 3 weeks ago

D is right  
upvoted 1 times

🗨️ 👤 **Yomzie** 6 months, 3 weeks ago

By using the Gremlin API, you can query a graph database in Azure Cosmos DB using the Gremlin query language. Gremlin is a graph traversal language and the primary query language for graph databases.

To query a graph database in Azure Cosmos DB using the Gremlin API, you would typically use a Gremlin console or a Gremlin client library. Here's a general outline of how you would query a graph database in Azure Cosmos DB using the Gremlin API:

Connect to the Cosmos DB account: You would need to provide the connection details for your Azure Cosmos DB account, including the URI and authentication keys or tokens.

Instantiate a Gremlin client: Use a Gremlin client library for your programming language of choice (e.g., Java, Python, .NET) to establish a connection to the Cosmos DB account.

Execute Gremlin queries: Use the methods provided by the Gremlin client library to execute Gremlin queries against the graph database in Azure Cosmos DB. These queries can traverse the graph, retrieve vertices (nodes), edges (relationships), and perform various graph operations. Process the results returned by the Gremlin queries as needed for your application logic.

upvoted 1 times

🗨️ 👤 **Goouter** 7 months, 3 weeks ago

Had this question today (01/26/2024). 'A' was no longer an option. I believe the correct answer is D.

upvoted 3 times

🗨️ 👤 **ikramus** 7 months, 2 weeks ago

and for the other question are those the same ?

upvoted 1 times

🗨️ 👤 **Kemsmems** 10 months, 1 week ago

" For example, you can use a document collection to store graph data side by side with documents, and use both SQL queries over JSON and Gremlin queries to query the collection." Answer is correct.

upvoted 2 times

🗨️ 👤 **Arshxd** 10 months, 3 weeks ago

Bing says: As nodes and edges by using the Gremlin language.

upvoted 1 times

🗨️ 👤 **Lou\_** 11 months, 3 weeks ago

the correct answer is D

upvoted 1 times

🗨️ 👤 **Prince9598** 1 year ago

D is correct

upvoted 1 times

🗨️ 👤 **rmeng** 1 year ago

You can query a graph database in Azure Cosmos DB as nodes and edges by using the Gremlin Language. Gremlin is a popular query language for graph databases, and it allows you to traverse and manipulate the graph structure of your data.

upvoted 1 times

🗨️ 👤 **Kraken84** 1 year, 1 month ago

I love how everyone is so "As a Matter of Fact" about "their" answer. Confidence can be a blessing...

upvoted 1 times

🗨️ 👤 **akhil5432** 1 year, 1 month ago

A is correct answer

upvoted 1 times

🗨️ 👤 **XtraWest** 1 year, 6 months ago

As per Chat GPT: Azure Cosmos DB supports querying graph databases using the Gremlin API. Gremlin is a graph traversal language that allows you to traverse and manipulate a graph database, and it is supported by Azure Cosmos DB. You can use the Gremlin API to perform graph queries, create vertices and edges, and traverse the graph to find paths and relationships between different nodes. The Gremlin API supports multiple programming languages, including Java, Python, and .NET, and it is widely used in graph database applications.

upvoted 3 times

🗨️ 👤 **chiiiweiii** 1 year, 6 months ago

Correct Answer Is D - As nodes and edges by using the Gremlin language.

upvoted 2 times

🗨️ 👤 **zerone72** 1 year ago

Please read carefully the linked article before suggesting wrong answers...  
upvoted 1 times

🗨️ 👤 **facilitator** 1 year, 6 months ago

wrong...Gremlin

The Gremlin API is used for graph databases. The MongoDB API stores data in the BSON format. The Table API is used to retrieve key/value pairs. The Cassandra API is used to retrieve tabular data.

upvoted 2 times

🗨️ 👤 **Joanna0** 1 year, 7 months ago

Gremlin!!!

<https://learn.microsoft.com/en-us/azure/cosmos-db/gremlin/tutorial-query>

upvoted 3 times

🗨️ 👤 **OrangeSG** 1 year, 9 months ago

Answer is D

Keyword is 'graph database', so it is Gremlin API.

The API for Gremlin is built based on Apache TinkerPop, a graph computing framework that uses the Gremlin query language.

Reference

Introduction to Azure Cosmos DB for Apache Gremlin

<https://learn.microsoft.com/en-us/azure/cosmos-db/gremlin/introduction>

upvoted 3 times

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

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

**Suggested Answer: B**

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.

Reference:

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

Community vote distribution

B (100%)

AVP\_Riga Highly Voted 3 years, 3 months ago  
AZ-900 is it you?  
upvoted 54 times

toshiro\_hitsugaya 1 year, 2 months ago  
This legit made me laugh. Good one  
upvoted 2 times

hmAZtime Highly Voted 3 years, 5 months ago  
right answer  
upvoted 11 times

AGTraining Most Recent 6 months, 2 weeks ago  
option B  
upvoted 1 times

AZFabio 1 year ago  
Selected Answer: B  
redundancy WITHIN a region = availability zones  
upvoted 1 times

akhil5432 1 year, 1 month ago  
option b  
upvoted 1 times

chiiiweiii 1 year, 6 months ago  
Selected Answer: B  
Correct Answer  
upvoted 4 times

Joanna0 1 year, 7 months ago  
Selected Answer: B  
Enabling redundancy across availability zones makes Azure Cosmos DB resilient to zone outages at the cost of increased charges.

<https://learn.microsoft.com/en-us/azure/cosmos-db/high-availability#replica-outages>  
upvoted 2 times

Mern 2 years, 3 months ago  
Selected Answer: B  
B is correct  
upvoted 2 times

🗨️ 👤 **GetulioJr** 3 years, 3 months ago

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. Availability Zones provide a 99.995% availability SLA with no changes to latency. In the event of a single zone failure, zone redundancy provides full data durability with RPO=0 and availability with RTO=0.

upvoted 6 times

🗨️ 👤 **chukismit** 3 years, 3 months ago

D righ answer

It is strongly recommended that you configure the Azure Cosmos accounts used for production workloads to enable automatic failover. This enables Cosmos DB to failover the account databases to available regions automatically. In the absence of this configuration, the account will experience loss of write availability for all the duration of the write region outage, as manual failover will not succeed due to lack of region connectivity.

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

upvoted 2 times

🗨️ 👤 **sri0509** 3 years, 3 months ago

B is the correct answer. The question is providing redundancy within a region and not across regions.

upvoted 3 times

🗨️ 👤 **JayBee65** 2 years, 9 months ago

The question states "... within an Azure region". The feature use describe operated across different regions NOT within the the same region, so D is wrong.

upvoted 1 times

What is a benefit of the Azure Cosmos DB Table API as compared to Azure Table storage?

- A. provides resiliency if an Azure region fails
- B. supports partitioning
- C. provides a higher storage capacity
- D. supports a multi-master model

**Suggested Answer: D**

Multi-master support for Azure Cosmos DB is now available in all public regions.

Azure CosmosDB table API is a key-value storage hosted in the cloud. It's a part of Azure Cosmos DB, that is Microsoft's multi-model database.

Reference:

<https://docs.microsoft.com/en-us/azure/cosmos-db/table-support>

Community vote distribution

D (100%)

  **keySEE** Highly Voted 3 years, 2 months ago

Just passed the exam with a 1000 score. The multi-master is the correst answer  
upvoted 87 times

  **Gbolatiti** 2 years, 11 months ago

And all these questions are still 100% valid?  
upvoted 5 times

  **nalwd** 1 year, 7 months ago

No bro  
upvoted 3 times

  **Armanas** 1 year, 6 months ago

... why do you say "No"?  
upvoted 4 times

  **sheilawu** 1 year ago

lol hack why  
upvoted 1 times

  **[Removed]** Highly Voted 1 year, 9 months ago

To all of you who wonder what multi-master support actually is:

'Till now, all the distributed system architecture works like one Write data server and multiple Read data servers. Data are replicated to all data servers from Write data server.

With the multi-master feature, we can choose all data servers to act as Write server data will be replicated with other data servers. The big challenge of multi-master, avoid the data duplication in data replication.'

<https://www.c-sharpcorner.com/article/utilize-multi-master-support-in-azure-cosmos-db/>

upvoted 18 times

  **GaneshNats** Most Recent 1 year, 4 months ago

Selected Answer: D

Supports a multi master model  
upvoted 2 times

  **chiiiweiii** 1 year, 6 months ago

Selected Answer: D

Correct Answer  
upvoted 2 times

🗨️ 👤 **Joanna0** 1 year, 7 months ago

Selected Answer: D

<https://azure.microsoft.com/en-gb/updates/cosmos-db-multi-master-support-now-generally-available/>

upvoted 1 times

🗨️ 👤 **Dusica** 1 year, 9 months ago

perhaps multi-master = multi region writes ; check second table in this article: <https://learn.microsoft.com/en-us/azure/cosmos-db/high-availability>

But then it might as well be A?

upvoted 1 times

🗨️ 👤 **CDL** 2 years, 2 months ago

D correct. tested in Azure portal.

1. create account in Azure Cosmo DB account
2. under left hand menu bar, locate "Settings" --> "Replicate data globally" <-- click it
3. Turn on "Enable" --> "Multi-region writers. explanation shown in "!" mark :

Multi-master capability allows you to take advantage of the provisioned throughput for your databases and containers across the globe.

upvoted 3 times

🗨️ 👤 **Nhendy** 2 years, 1 month ago

it says "Multi-region capability allows..." not Multi-master capability !! why would you lie in something like that?

upvoted 1 times

🗨️ 👤 **jerkyflexoff** 3 years, 3 months ago

Looks like D: - New Feature -

<https://azure.microsoft.com/en-gb/updates/cosmos-db-multi-master-support-now-generally-available/>

upvoted 3 times

🗨️ 👤 **jerkyflexoff** 3 years, 2 months ago

This is also the answer in the MeasureUp questions

upvoted 1 times

🗨️ 👤 **JayRee** 3 years, 5 months ago

An overview of Azure Cosmos DB multi-master capabilities

<https://azure.microsoft.com/en-us/resources/videos/ignite-2018-multi-master-at-global-scale-with-azure-cosmos-db/>

upvoted 4 times

🗨️ 👤 **fflyin2k** 3 years, 5 months ago

A should be correct answer. (Agree with hmAZtime)

According to following link,

- 'master', 'model' are not mentioned at all.

- Global distribution is mentioned.

<https://docs.microsoft.com/en-us/azure/cosmos-db/table-support>

upvoted 8 times

🗨️ 👤 **hmAZtime** 3 years, 5 months ago

I think it's A Table API support failover over the region. From the same document provided: "Support for automatic and manual failovers at any time, anywhere in the world."

upvoted 2 times

🗨️ 👤 **hmAZtime** 3 years, 5 months ago

I think I would agree with the answer after one more time research. While not found any 'multi - master' proof in Microsoft documents, I did see couple of other sites mentioned 'multi - modle', I pasted one link here, sorry I lost the other one. on the other hand, A is not that clearly described to me, I'm not sure it meant regionally failover. So this time I go with D.

<https://db-engines.com/en/system/Microsoft+Azure+Cosmos+DB%3BMicrosoft+Azure+Table+Storage>

upvoted 2 times

🗨️ 👤 **zvasanth2** 3 years, 2 months ago

multi master

<https://azure.microsoft.com/en-gb/updates/cosmos-db-multi-master-support-now-generally-available/>

upvoted 2 times

Your company needs to design a database that shows how changes in network traffic in one area of a network affect network traffic in other areas of the network.

Which type of data store should you use?

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

**Suggested Answer: A**

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.

Reference:

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

Community vote distribution

A (100%)

- 🗨️ **hmAZtime** Highly Voted 3 years, 5 months ago  
right answer  
upvoted 22 times
- 🗨️ **lhueghian** Highly Voted 3 years, 3 months ago  
Graph: Taking about relationship between one area with the other, so it makes sense that it's Graph  
upvoted 17 times
- 🗨️ **AGTraining** Most Recent 6 months, 2 weeks ago  
option A  
upvoted 1 times
- 🗨️ **akhil5432** 1 year, 1 month ago  
Correct answer is A  
upvoted 1 times
- 🗨️ **mamahani** 1 year, 5 months ago  
I took the exam 12.04.2023; this question appeared, correct answer:  
graph  
upvoted 9 times
- 🗨️ **chiiweiii** 1 year, 6 months ago  
Selected Answer: A  
Correct Answer  
upvoted 3 times
- 🗨️ **bfan95** 2 years, 4 months ago  
Selected Answer: A  
Right answer  
upvoted 2 times
- 🗨️ **Mlynkowsky** 2 years, 5 months ago  
Selected Answer: A  
A is right  
upvoted 1 times

HOTSPOT -

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE: Each correct selection is worth one point.

Hot Area:

### Answer Area

Statements	Yes	No
Azure Databricks can consume data from Azure SQL Database	<input type="radio"/>	<input type="radio"/>
Azure Databricks can consume data from Azure Event Hubs	<input type="radio"/>	<input type="radio"/>
Azure Databricks can consume data from Azure Cosmos DB	<input type="radio"/>	<input type="radio"/>

Suggested Answer:

### Answer Area

Statements	Yes	No
Azure Databricks can consume data from Azure SQL Database	<input checked="" type="radio"/>	<input type="radio"/>
Azure Databricks can consume data from Azure Event Hubs	<input checked="" type="radio"/>	<input type="radio"/>
Azure Databricks can consume data from Azure Cosmos DB	<input checked="" type="radio"/>	<input type="radio"/>

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.

Reference:

<https://docs.microsoft.com/en-us/azure/databricks/data/data-sources/sql-databases-azure> <https://docs.microsoft.com/en-us/azure/databricks/scenarios/databricks-stream-from-eventhubs>

 **Alex210** Highly Voted 1 year, 10 months ago  
marketing level 100  
upvoted 27 times

 **AgileRoc** Highly Voted 3 years, 3 months ago  
correct answer

upvoted 20 times

🗨️ 👤 **AGTraining** Most Recent 6 months, 2 weeks ago

All 'Yes'

upvoted 1 times

🗨️ 👤 **[Removed]** 1 year, 4 months ago

Yes, correct.

Databricks is a DW tool to process and integrate data from all these sources.

upvoted 2 times

🗨️ 👤 **chiiiweiii** 1 year, 6 months ago

Correct Answer

upvoted 2 times

🗨️ 👤 **dayan911** 3 years, 2 months ago

All 'Yes'

upvoted 2 times

🗨️ 👤 **hibana2077** 3 years, 4 months ago

agree.

upvoted 2 times

🗨️ 👤 **hmAZtime** 3 years, 5 months ago

agree.

upvoted 3 times

🗨️ 👤 **Drummer** 3 years, 5 months ago

Correct answer

upvoted 2 times

DRAG DROP -

Match the datastore services to the appropriate descriptions.

To answer, drag the appropriate service from the column on the left to its description on the right. Each service may be used once, more than once, or not at all.

NOTE: Each correct match is worth one point.

Select and Place:

Services	Answer Area
Azure Blob storage	Service Enables the use of SQL queries against data stored in JSON documents
Azure Cosmos DB	Service Enables users to access data by using the Server Message Block (SMB) version 3 protocol
Azure Files	
Azure Table storage	

#### Suggested Answer:

Services	Answer Area
Azure Blob storage	Azure Cosmos DB Enables the use of SQL queries against data stored in JSON documents
Azure Cosmos DB	Azure Files Enables users to access data by using the Server Message Block (SMB) version 3 protocol
Azure Files	
Azure Table storage	

Box 1: Azure Cosmos DB -

In Azure Cosmos DB's SQL (Core) API, items are stored as JSON. The type system and expressions are restricted to deal only with JSON types.

Box 2: Azure Files -

Azure Files offers native cloud file sharing services based on the SMB protocol.

Reference:

<https://docs.microsoft.com/en-us/azure/cosmos-db/sql-query-working-with-json> <https://cloud.netapp.com/blog/azure-smb-server-message-block-in-the-cloud-for-azure-files>

-  **GeoTan** Highly Voted 3 years, 3 months ago  
 answer correct  
 upvoted 20 times
-  **hmAZtime** Highly Voted 3 years, 5 months ago  
 right answer  
 upvoted 10 times
-  **AGTraining** Most Recent 6 months, 2 weeks ago  
 Azure cosmos DB  
 Aure Files  
 upvoted 1 times
-  **Abhifzd** 1 year, 4 months ago  
 answer correct  
 upvoted 2 times
-  **Kseetha** 1 year, 6 months ago

Azure cosmos DB

Azure Files

upvoted 4 times

  **chiiweiii** 1 year, 6 months ago

Correct Answer

upvoted 3 times

  **CDL** 2 years, 2 months ago

correct

for "Azure File" ref: <https://docs.microsoft.com/en-us/azure/storage/files/storage-how-to-use-files-windows>

upvoted 3 times

  **Drummer** 3 years, 5 months ago

The answer is correct

upvoted 5 times

HOTSPOT -

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

Hot Area:

### Answer Area

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

	▼
account	
container	
database	
item	

Suggested Answer:

### Answer Area

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

	▼
account	
container	
database	
item	

Reference:

<https://docs.microsoft.com/en-us/azure/cosmos-db/create-cosmosdb-resources-portal>

- 🗨️ **Rakesh16** Highly Voted 2 years, 11 months ago  
Account is correct.  
upvoted 20 times
- 🗨️ **AGTraining** Most Recent 6 months, 2 weeks ago  
Option A  
upvoted 1 times
- 🗨️ **scou** 11 months, 3 weeks ago  
Appeared on 30/09/23 exam  
upvoted 1 times
- 🗨️ **akhil5432** 1 year, 1 month ago  
Correct answer  
upvoted 2 times
- 🗨️ **VNR25** 1 year, 2 months ago  
This question appeared in the 07/07/23 exam and the answer is correct.  
upvoted 2 times
- 🗨️ **Abhifzd** 1 year, 4 months ago  
Account is correct.  
upvoted 3 times
- 🗨️ **chiiweiii** 1 year, 6 months ago  
Correct Answer  
upvoted 2 times
- 🗨️ **reza88** 1 year, 11 months ago  
it was a Q 20/10/2022  
upvoted 4 times
- 🗨️ **RichMelbourne** 2 years, 5 months ago  
Shown on test 15/04/2022  
upvoted 4 times

