

AWS BDS-C00 Exam

Volume: 297 Questions

Question: 1

In regard to DynamoDB, when you create a table with a hash-and-range key _____

- A. You can optionally define one or more secondary indexes on that table
- B. You must define one or more secondary indexes on that table
- C. You must define one or more Global secondary indexes on that table
- D. You must define one or more Local secondary indexes on that table

Answer: A

Explanation:

When you create a table with a hash-and-range key, you can optionally define one or more secondary indexes on that table. A secondary index lets you query the data in the table using an alternate key, in addition to queries against the primary key.

Reference:

<http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/DataModel.html>

Question: 2

Amazon DynamoDB supports these scalar data types: _____.

- A. Number and String
- B. Number and Binary
- C. Number, String, and Binary
- D. Number, String, Binary and Datetime

Answer: C

Explanation:

Amazon DynamoDB supports three scalar data types: Number, String, and Binary. Additionally, Amazon DynamoDB supports multi-valued types: Number Set, String Set, and Binary Set.

Reference:

<http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/SecondaryIndexes.html>

AWS BDS-COO Exam

Question: 3

True or false: In DynamoDB, it is up to you to manage the partitioning and re-partitioning of your data over additional DynamoDB tables if you need additional scale.

- A. True, It is optional to re-partition by yourself or automatically.
- B. False, DynamoDB automatically partitions your data and workload.
- C. False, the table size is fixed and you cannot re-partition it.
- D. True, AWS DynamoDB does automatic partitioning and SSD technologies to meet your throughput requirements and deliver low latencies at any scale.

Answer: B

Explanation:

Amazon DynamoDB automatically partitions your data and workload over a sufficient number of servers to meet the scale requirements you provide.

Reference:

<https://aws.amazon.com/dynamodb/faqs/>

Question: 4

Complete this statement: "When you load your table directly from an Amazon _____ table, you have the option to control the amount of provisioned throughput you consume."

- A. DataPipeline
- B. S3
- C. DynamoDB
- D. RDS

Answer: C

Explanation:

When you load your table directly from an Amazon DynamoDB table, you have the option to control the amount of Amazon DynamoDB provisioned throughput you consume.

Reference:

http://docs.aws.amazon.com/redshift/latest/dg/t_Loading_tables_with_the_COPY_command.html

AWS BDS-C00 Exam

Question: 5

In DynamoDB, which of the following operations is not possible by the console?

- A. Copying an item
- B. Updating an item
- C. Deleting an item
- D. Blocking an item

Answer: D

Explanation:

By using the console to manage DynamoDB, you can perform the following: adding an item, deleting an item, updating an item, and copying an item.

Reference:

<http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/AddUpdateDeleteItems.html>

Question: 6

DynamoDB uses _____ only as a transport protocol, not as a storage format.

- A. JSON
- B. XML
- C. SGML
- D. WDDX

Answer: A

Explanation:

DynamoDB uses JSON only as a transport protocol, not as a storage format. The AWS SDKs use JSON to send data to DynamoDB, and DynamoDB responds with JSON, but DynamoDB does not store data persistently in JSON format.

Reference:

<http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Programming.LowLevelAPI.html>

AWS BDS-C00 Exam

Question: 7

When you create a table in DynamoDB, which one of the following bits of information is not obligatory to be provided?

- A. Units of Capacity required for reads
- B. Range key
- C. Hash key
- D. Units of Capacity required for writes

Answer: B

Explanation:

To create a table in DynamoDB, you should provide the table name, the attribute name for the primary Hash key, as well as throughput units, read and write units capacity.

Reference:

<http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/GettingStartedCreateTables.html>

Question: 8

In the DynamoDB console, you can choose the ___ tab to view some key CloudWatch metrics for a selected table.

- A. Browse Items
- B. Details
- C. Alarms
- D. Metrics

Answer: D

Explanation:

In the DynamoDB console, you can choose the Metrics tab to view some key CloudWatch metrics for a selected table.

Reference:

http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/MonitoringConsole_DynamoDB.html

AWS BDS-COO Exam

Question: 9

Which of the following statements is correct of Amazon DynamoDB?

- A. Data in DynamoDB cannot be shifted to Amazon Redshift. Instead, data can be shifted to Amazon CloudWatch.
- B. Every item in a DynamoDB table is identified by a foreign key, which allows you to quickly access data items.
- C. DynamoDB does not support multiple native data types (numbers, strings, binaries, and multi-valued attributes).
- D. DynamoDB does not have a fixed schema. Instead, each data item may have a different number of attributes.

Answer: D

Explanation:

DynamoDB does not have a fixed schema. Instead, each data item may have a different number of attributes. Multiple data types (strings, numbers, binary, and sets) add richness to the data model.

Reference: <http://awsdocs.s3.amazonaws.com/dynamodb/latest/dynamodb-dg.pdf>

Question: 10

What kind of service is provided by AWS DynamoDB?

- A. Relational Database
- B. Document Database
- C. Dynamic Database
- D. NoSQL Database

Answer: D

Explanation:

DynamoDB is a fast, fully managed NoSQL database service.

Reference: <http://aws.amazon.com/dynamodb/>

Question: 11

AWS BDS-COO Exam

True or false: For an online game, it is better to use DynamoDB than a relational database engine.

- A. This is true because DynamoDB fully manages your database and can grow with your application requirements.
- B. This is false.
- C. This is true because DynamoDB has all the functionalities of a relational database.
- D. This is false because a relational database engine on Redshift has more functionalities than DynamoDB.

Answer: A

Explanation:

This is true. An online game might start out with only a few thousand users and a light database workload consisting of 10 writes per second and 50 reads per second. However, if the game becomes successful, it may rapidly grow to millions of users and generate tens (or even hundreds) of thousands of writes and reads per second. It may also create terabytes or more of data per day. Developing your applications against Amazon DynamoDB enables you to start small and simply dial-up your request capacity for a table as your requirements scale, without incurring downtime. Amazon DynamoDB gives you the peace of mind that your database is fully managed and can grow with your application requirements.

Reference: <http://aws.amazon.com/dynamodb/faqs/>

Question: 12

In regard to DynamoDB, for which one of the following parameters does Amazon not charge you?

- A. I/O usage within the same Region
- B. Cost per provisioned write units
- C. Storage cost
- D. Cost per provisioned read units

Answer: A

Explanation:

In DynamoDB, you will be charged for the storage and the throughput you use rather than for the I/O which has been used.

AWS BDS-COO Exam

Reference: <http://aws.amazon.com/dynamodb/pricing/>

Question: 13

Which of the following does Amazon DynamoDB perform?

- A. Neither increment nor decrement operations
- B. Only increment on vector values
- C. Atomic increment or decrement on scalar values
- D. Only atomic decrement operations

Answer: C

Explanation:

Amazon DynamoDB allows atomic increment and decrement operations on scalar values.

Reference: <http://aws.amazon.com/dynamodb/faqs/>

Question: 14

A gaming company comes to you and asks you to build them infrastructure for their site. They are not sure how big they will be as with all start ups they have limited money and big ideas. What they do tell you is that if the game becomes successful, like one of their previous games, it may rapidly grow to millions of users and generate tens (or even hundreds) of thousands of writes and reads per second. After considering all of this, you decide that they need a fully managed NoSQL database service that provides fast and predictable performance with seamless scalability. Which of the following databases do you think would best fit their needs?

- A. Amazon DynamoDB
- B. Amazon SimpleDB
- C. Amazon Redshift
- D. Any non-relational database.

Answer: A

Explanation:

Amazon DynamoDB is a fully managed NoSQL database service that provides fast and predictable performance with seamless scalability. Amazon DynamoDB enables customers to

AWS BDS-COO Exam

offload the administrative burdens of operating and scaling distributed databases to AWS, so they don't have to worry about hardware provisioning, setup and configuration, replication, software patching, or cluster scaling.

Today's web-based applications generate and consume massive amounts of data. For example, an online game might start out with only a few thousand users and a light database workload consisting of 10 writes per second and 50 reads per second. However, if the game becomes successful, it may rapidly grow to millions of users and generate tens (or even hundreds) of thousands of writes and reads per second. It may also create terabytes or more of data per day. Developing your applications against Amazon DynamoDB enables you to start small and simply dial-up your request capacity for a table as your requirements scale, without incurring downtime. You pay highly cost-efficient rates for the request capacity you provision, and let Amazon DynamoDB do the work over partitioning your data and traffic over sufficient server capacity to meet your needs. Amazon DynamoDB does the database management and administration, and you simply store and request your data. Automatic replication and failover provides built-in fault tolerance, high availability, and data durability. Amazon DynamoDB gives you the peace of mind that your database is fully managed and can grow with your application requirements.

Reference: <http://aws.amazon.com/dynamodb/faqs/>

Question: 15

Could you use IAM to grant access to Amazon DynamoDB resources and API actions?

- A. Yes
- B. Depended to the type of access
- C. No
- D. In DynamoDB there is no need to grant access

Answer: A

Explanation:

Amazon DynamoDB integrates with AWS Identity and Access Management (IAM). You can use AWS IAM to grant access to Amazon DynamoDB resources and API actions. To do this, you first write an AWS IAM policy, which is a document that explicitly lists the permissions you want to grant. You then attach that policy to an AWS IAM user or role.

Reference:

<http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/UsingIAMWithDDB.html>

AWS BDS-COO Exam

Question: 16

Which of the statements below best define Amazon Kinesis Streams? - (choose three)

- A. Each record in the stream has a sequenced number that is assigned by Kinesis Streams.
- B. An ordered sequence of data records
- C. A record is the unit of data stored in the Amazon Kinesis Stream
- D. A random sequence of data records

Answer: ABC

Explanation:

An Amazon Kinesis Stream is an ordered sequence of data records. Each record in the stream has a sequenced number that is assigned by Kinesis Streams.

Reference: <http://docs.aws.amazon.com/streams/latest/dev/working-with-kinesis.html>

Question: 17

Which of the following statements about Amazon Kinesis streams are correct? (choose three)

- A. Kinesis Streams can support up to 1,000 records per second for writes and up to a maximum total data write rate of one megabyte per second, including partition keys.
- B. You cannot increase or decrease the number of shards in a stream.
- C. Each shard can support up to five transactions per second for reads
- D. Kinesis streams can support up to a maximum total data read rate of two megabytes per second.

Answer: ACD

Question: 18

Which network connections are used by AWS Snowball to minimize data transfer times?

- A. both thinnet and thicknet copper cables
- B. both RJ45 and SFP+ with either a fiber or copper interface
- C. both USB and Ethernet cables with special adapters

AWS BDS-COO Exam

D. both UTP and STP copper cables

Answer: B

Explanation:

An AWS Snowball appliance has a 10GBaseT network connection (both RJ45 as well as SFP+ with either a fiber or copper interface) to minimize data transfer times.

Reference: <https://aws.amazon.com/snowball/details/>

Question: 19

You want to export objects in an S3 bucket using the AWS Snowball Management Console. Assume you have a bucket containing the following objects, sorted in UTF-8 binary order:

01

Aardvark

Aardwolf

Aasvogel/apple

Aasvogel/banana

Aasvogel/cherry

Banana

Car

What happens if you specify no beginning range and set Aasvogel as the ending range?

A. No objects are exported; you must specify a beginning range.

B. Objects Aasvogel/apple, Aasvogel/banana and Aasvogel/cherry are exported.

C. Objects 01, Aardvark, and Aardwolf are exported.

D. Objects 01, Aardvark, Aardwolf, Aasvogel/apple, Aasvogel/banana and Aasvogel/cherry are exported.

Answer: D

Explanation:

When you create an export job in the AWS Snowball Management Console, you can choose to export an entire Amazon S3 bucket or a specific range of objects keys. To export a range, you define the length of the range by providing either an inclusive range beginning, an inclusive range ending, or both. Ranges are UTF-8 binary sorted.

Reference: <http://docs.aws.amazon.com/snowball/latest/ug/using-console.html>

Question: 20