

AWS Certified Big Data - Specialty

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(s): A

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(s): C

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(s): B

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(s): C

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(s): D

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

A. JSON

B. XML

C. SGML

D. WDDX

Answer(s): A

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(s): B

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(s): D

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(s): D

10. What kind of service is provided by AWS DynamoDB?

A. Relational Database

B. Document Database

C. Dynamic Database

D. NoSQL Database

Answer(s): D

11. 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(s): A

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(s): A

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(s): C

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(s): A

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(s): A

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

B. A record is the unit of data stored in the Amazon Kinesis Stream

C. A random sequence of data records

Answer(s): A B

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(s): A C D

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

D. both UTP and STP copper cables

Answer(s): B

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(s): D

20. Sources of data that can be imported with the Snowball client are: (1) Files or objects hosted in locally mounted file systems, and (2) files or objects from a(n) .

A. HDFS cluster

B. NFS server

C. AFS server

D. GPFS cluster

Answer(s): A
