Amazon AWS-BigData-Specialty Exam

Volume: 80 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

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

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

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

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