

By OnlineInterviewQuestions.com

The use of SQL databases started plummeting with the advent of DynamoDB. The DynamoDB is pretty much being used as the replacement of SQL database. This replacement was made due to the key advantage which is, the DynamoDB is extremely fast and is highly scalable.

With the increase of use of this [NoSQL database](#) service, the industry demands candidates with DynamoDB knowledge. It is high time for the aspirants who seek a bright future in the back-end field to be well-aware of this database. And, if you already are acknowledged with DynamoDB and your interview is heading near, then this list of few **DynamoDB Interview Questions** will help you crack the interview.

Q1. Explain what is DynamoDB?

It is a NoSQL database service which is superintended and provides fast and inevitable performance with seamless scalability.

It also enables the users to offload the operating and scaling distributed databases to AWS.

So that the users don't have to worry about setup and configuration, hardware provisions, replication, throughput capacity planning, cluster scaling or software patching.

Q2. List some benefits of using Amazon DynamoDB?

Here is the list of some of the benefits of using Amazon DynamoDB:

- It is a managed service where there is no need to hire the experts or worry about installation, setup, cluster etc.,
- It is scalable.
- It provides the users high throughput at very low latency.
- It is durable and highly available.
- It is flexible and allows dynamic tables creation that includes multi-valued attributes.
- It is cost-effective.

Q3. What are the non-relational Databases?

The Non-Relational databases are NoSQL databases.

These databases are categorized into four groups, and they are:

- Key-value stores
- Graph stores
- Column stores
- Document stores

Q4. Is DynamoDB free for use?

In Amazon DynamoDB you only pay for the resources you provision.

Get started with free tier limits of DynamoDB on which many applications operate. In need of resources, the pricing per month varies depending on which type of resources you require.

Q5. Explain what is DynamoDB Mapper Class?

It is the entry point to DynamoDB. The DynamoDB Mapper class provides access to a DynamoDB endpoint and enables the user to access to their data in various tables, execute queries and scan against tables, and perform CRUD operations on items.

Q6. List the Data Types supported by DynamoDB?

DynamoDB supports four scalar data types, and they are:

- Number
- String
- Binary
- Boolean

DynamoDB supports collection data types such as:

- Number Set
- String Set
- Binary Set
- Heterogeneous List
- Heterogeneous Map

DynamoDB also supports Null values.

Q7. Does DynamoDB support in-place atomic updates?

Amazon DynamoDB supports quick in-place atomic updates, where the numeric attribute's increment and decrement can be done in arrow using just one API call or similarly, you can add or remove sets, lists, or maps.

Q8. What kind of query functionality does DynamoDB support?

- It supports GET/PUT operation using the user-defined primary key.
- It provides flexible querying by letting query a non-primary key attribute using local secondary indexes and Global secondary indexes.
- It allows quick reads and writes data for an item associated with single attribute partition primary key.
- It allows you to use the Query API to retrieve all the items for a single composite partition-sort key across a range of sort keys.

Q9. Does Amazon DynamoDB support conditional operations?

For an operation to be completed on an item, you have to specify a condition. You can define a ConditionExpression that can be constructed from the following:

- Boolean functions: ATTRIBUTE_EXIST, CONTAINS, and BEGINS_WITH
- Comparison operators: =, <>, <, >, <=, >=, BETWEEN, and IN
- Logical operators: NOT, AND, and OR.

You can also construct a free-form conditional expression that combines multiple conditional clauses which also includes nested clauses.

Q10. List some difference between Amazon DynamoDB and Amazon SimpleDB?

- **Amazon DynamoDB :** It is a fast and scalable NoSQL Database Service that is designed for internet scale applications, to maintain predictably high performance and to be highly cost-effective for workloads of any scale, on the top of all it is highly recommended.
- **Amazon SimpleDB :** It has scaling limitations and is a good fit for smaller workloads that requires query flexibility. It automatically indexes all item attributes and supports query flexibility at the cost of performance and scale.

Q11. List the APIs provided by Amazon DynamoDB?

- CreateTable
- UpdateTable
- DeleteTable
- DescribeTable
- ListTables
- PutItem
- BatchWriteItem
- UpdateItem
- DeleteItem
- GetItem
- BatchGetItem.
- Query
- Scan

Q12. What Amazon DynamoDB BatchWriteItem API does?

The Amazon DynamoDB BatchWriteItem inserts, deletes and replaces multiple items across multiple tables in just one request, but not in one transaction.

With maximum total request size of 16 MB, it supports batches of up to 25 items to insert or delete.

Q13. Please explain key-value store?

It is a database service that provides and supports storing, updating and querying the objects that are identified using key and values that constitutes the actual content that is being stored.

Q14. What is DynamoDB Auto Scaling?

DynamoDB is a superintended feature that scales up and down the read and write capacity of a DynamoDB or a global secondary index automatically.

Q15. Please explain what are global secondary indexes?

An index with a different partition and partition-and-sort key from those on the table is called global Secondary index.

It is considered to be “global” in the sense that queries on the index can span all the items in a table across, all the partitions.

Q16. List types of secondary indexes supported by Amazons DynamoDB?

Two types of indexes are supported by Amazon DynamoDB. They are:

- Global Secondary index – It is an index with a partition or a partition sort key that is different from those on the table. The global secondary index is considered to be global because queries on the index can span all the items in a table, across all the partitions.
- Local secondary index – An index that has the same partition key as that of the table but different sort key. It is considered to be “local” because every partition of the index is scoped to a table partition that has the same partition key.

Q17. How many numbers of global secondary indexes do you create per table?

You are allowed to create a maximum of 5 global secondary indexes per table.

Q18. What are API calls supported by a global secondary index?

“Query” and “Scan” are the API calls that are supported by Global Secondary Index.

Q19. How to delete a Global Secondary Index in Amazon DynamoDB?

A Global secondary index can be deleted from the console or through an API call.

Select the table from which you want to delete the Global Secondary index on the console, select the “indexes” tab under “Table items” and then click on the “Delete” button which is next to the delete index.

A Global Secondary Index can also be deleted using the Update Table API call.

Q20. What are local secondary indexes?

An index with the same partition key as the table but different sort key is called local secondary index.

A local secondary index is considered to be “local” because every partition of a local secondary index is scaled to a table partition which contains the same partition key.

Q21. What are Projections?

Projections are the set of attributes that are copied or projected from the table to an index.

They are in addition to the index key attributes and primary key attributes which are automatically projected.

Whenever you define a local secondary index, it is required to specify the attributes that are projected into the index.

Each index consists of a minimum of 3 attributes, and they are:

- Table partition key value
- Attribute to serve as the index sort key
- Table sort key value

Q22. How to create local secondary indexes?

A Local secondary index needs to be created at the time of creating a table that can't be added currently.

So, later when you create a local secondary index, try specifying the following two parameters.

- **Indexed Sort Key** : The attribute on which indexing and querying are performed.
- **Projected Attributes** : It is the list of attributes from the table that can directly be copied into the local secondary index. Without projected attributes, the local secondary index contains only primary keys and secondary keys.

Q23. Can I add local secondary indexes to an existing table?

When you create a table with local secondary indexes, you can create a local secondary index to use it in future by defining a sort key element which is not used currently.

So, at this time it is probably not possible to add local secondary indexes in the existing table.

Q24. How many local secondary indexes can be created on one table?

Each table can constitute up to five local secondary indexes.

Q25. Can I delete local secondary indexes?

Currently, Amazon DynamoDB is incapable of removing the local secondary indexes from the table once they are created, but it is possible to delete the entire table.

Please Visit OnlineInterviewquestions.com to download more pdfs