By OnlineInterviewQuestions.com

MongoDB Interview Questions

MongoDB Interview Questions and Answers

Q1. What do you know about MongoDB?

MongoDB is a cross-platform document-oriented database program that is open source and free in nature. It can also be classified as the NoSQL database program. It was developed by MongoDB Inc. JSON-like documents with schema are used by MongoDB. Field, range queries, and regular expressions are also supported by MongoDB. The queries used by this contain user-defined JavaScript functions. It provides high availability with the replica sets. Fields can be indexed with primary and secondary indices. Sharding is used by MongoDB for scaling horizontally which determines how the data will be distributed. It can also be used as a file system with load balancing. It can also be used for batch processing of data.

Q2. List the important features of MongoDB.

The important features of MongoDB are listed below: –

1. **Aggregation framework**– it uses aggregation framework for the batch processing of data and aggregation operations. It should be used in an effective manner.
2. **The usage of BSON format in MongoDB**– it uses the binary-encoded serialization of JSON-like documents. The data-types like- date and binary are supported in the format.
3. **The sharing feature of MongoDB**– using this feature, MongoDB supports the distribution of data across multiple machines. It supports deployment with large sets of data.
4. **The Ad hoc queries in MongoDB**– it supports field, range queries, and regular expressions.
5. **Collections**– MongoDB supports fixed size collection which can also be termed as capped collections.

Q3. Which all languages can be used with MongoDB?

Here goes a list of the languages which can be used with MongoDB: –

- C
- C++
- C#
- Java
- Node.js
- Perl

- PHP
- Python
- Ruby
- Scala
- Go
- Erlang.

Currently, MongoDB provides driver support for the languages listed above. It can be used easily with any of these languages.

Q4. What is the use of a namespace in MongoDB?

In MongoDB, a namespace is a combination of the database name and collection or index name. It is a canonical name for an index or a collection in MongoDB. A namespace consists of all the documents in MongoDB. The maximum length of a collection of namespaces is called the namespace length. It consists of the database name, a dot operator (.), and the name of the collection.

Syntax- <database>.<collection>

The namespace contains all the helper classes which can be used to construct various options which are to be used in the drivers. A namespace can also be termed as the concatenation of the collection name and the database name.

Q5. What do you mean by a replica set in MongoDB?

A group of mongo instances which is able to host the same data set is known as a replica set in MongoDB. A replica set consists of a primary node and a secondary node too. With the help of a replica set, all the data from primary node to the secondary node replicates. Replication is a process of synchronizing the data. Replication provides redundancy and it also increases the availability of data with the help of multiple copies of data on the different database server. It also protects the database from the loss of a single server.

Q6. What should all points be taken into consideration while creating a schema in MongoDB?

Following are the points which should be taken into consideration when you create a schema in MongoDB: –

- Your schema should be designed according to your requirements.
- If the objects are to be used together, you should combine them into one document. Otherwise, you should use them separately.
- Joins should be performed while writing and not while reading.
- If use cases are used more frequently, then you should optimize your schema.
- Complex aggregation should be done in the schema.

Q7. What do you understand by a profiler with reference to MongoDB?

A database profiler is something which collects fine-grained data about write operations of MongoDB, the cursors, and the database commands. Profiling can be enabled on a per database or pre instance basis. When profiling is enabled, the profiling level also gets configurable. By default, the profiler will be off. All the data collected by the system.profiler collection is written by the database profiler to it. This collection is a capped collection. 0, 1, and 2, these are the three levels available in the database profiler. Profiling in a database can be enabled from the mongo shell or using a profile command through the driver. While enabling profiling, profiling level is also set. The profiler is used to record data in the system.profile collection.

Q8. Write the syntax for creating a collection and dropping a collection in MongoDB.

The creation of the collection is the basic step in MongoDB. Here is the syntax for creating a collection in MongoDB-

```
 Db.createCollection (name, options)
```

Dropping a collection basically means deleting a collection. Following is the syntax for dropping a collection in MongoDB-

```
Db.collection.drop()
```

Q9. What are the different types of NoSQL databases? Give some examples.

There are 4 basic types of NoSQL database. They are as follows: –

- Key value store NoSQL database
- Document store NoSQL database
- Column store NoSQL database
- Graph-based NoSQL database

The various examples of NoSQL database are- MongoDB, Cassandra, CouchDB, Hypertable, Redis, Riak, Neo4j, HBase, Couchbase, MemcacheDB, Voldemort, RevenDB, etc.

Q10. What are the differences between MongoDB and MySQL?

The various differences between MongoDB and MySQL are listed below:

- **In terms of data representation**– in MySQL, we represent data in the form of tables and rows. Whereas, in MongoDB data is represented as collections of JSON documents.
- **Querying**– in SQL, we put together a string in the query language which is then parsed by the database system. Whereas, in MongoDB, object querying is used.
- **Relationships**– in MySQL, the relational database is only the Join operation which allows us to perform queries across multiple tables. Whereas MongoDB does not support join operation but can support multi-dimensional data types such as arrays, etc. in this, embedding is a process in which we place one document inside the other.
- **Transactions**– MySQL supports atomic transactions which are the ability to contain multiple operations within a transaction. Whereas, MongoDB does not support transactions.
- **Schema definition**– in MySQL, you need to define your tables and columns before storing anything. Whereas, in MongoDB, you don't need to define a schema.

Q11. Explain sharding in MongoDB?

Q12. List some alternatives to MongoDB?

Q13. I am getting error cannot overwrite model once compiled Mongoose, How to fix it?

Q14. What is the "__v" field in Mongoose?

When the mongoose is created at the first time the version key is a property set on every document. The value of this key comprises of the internal revision of the document. It is understood that the name of this document is configurable. The default key is __v

Q15. Explain what is MongoDB?

MongoDB is an open-source NoSQL database management program that supports various forms of data. MongoDB is basically used as a tool that can manage document-oriented information, store or retrieve information.

Q16. List some important features of MongoDB?

Some of the important features of MongoDB are as follows:

- Support ad hoc queries.
- Indexing and Replication.
- Duplication of data.
- Load balancing.
- Supports map reduction and aggregation tools.
- Uses JavaScript instead of Procedures.
- It is a schema-less database written in C++.

Q17. What is namespace in MongoDB?

A **namespace** is a canonical name for a collection or index in MongoDB. The namespace is basically a combination of the database name and the name of the collection or index, like so: [database-name]. [collection-or-index-name] .

Q18. What is BSON in MongoDB?

BSON stands for Binary JSON which is a binary serialization of JSON-like documents that is used by MongoDB to store the documents in a BSON format. BSON implementation is used for supporting embedding objects and arrays within other objects.

Q19. What type of DBMS is MongoDB?

MongoDB is an open-source **NoSQL** type DBMS or database management program that supports various forms of data.

Q20. What is the document structure of MongoDB?

MongoDB uses **JSON format** as the structure of the document. JSON documents support embedded fields, so related data and lists of data can be stored with the document instead of an external table.

Q21. What is replica set in MongoDB?

In MongoDB, a **Replica Set** is used to carry out MongoDB replication, which is a process of creating a copy of the same data set in more than one MongoDB server. A replica set is basically a group of MongoDB instances that maintain the same data set and pertain to any MongoDB process.

Q22. What is profiler in MongoDB?

In MongoDB, the database **profiler** collects detailed information about Database Commands executed against a running MongoDB instance. It writes all the data it collects to a system and captures and records data on the performance of write operations, cursors, and database commands on a running MongoDB instance.

Q23. Write the syntax for creating and droping a collection in MongoDB?

The syntax for creating and droping a collection in MongoDB are as follows:

- **Creating a Collection: collection_name. insert({key:value, key:value…})**
- **Droping a Collection: db.collection.drop({ writeConcern: })**

Q24. What is the size limit of a document?

The maximum Document Size Limit in MongoDb is **16 megabytes** or 16777216 bytes. The maximum document size helps ensure that a single document cannot use an excessive amount of RAM or, during transmission, an excessive amount of bandwidth.

Q25. What is _id Field in MongoDB?

In MongoDB, **_id field** is always the first field in the documents therefore each document stored in a collection requires a unique _id field that acts as a primary key.

Q26. Explain what is ObjectId in MongoDB?

An **ObjectId** in MongoDB is a 12-byte BSON type in which the first 4 bytes of the ObjectId represent the time in seconds then the next 3 bytes of the ObjectId represent the machine identifier then again the next 2 bytes are of process id and the last Field is 3 bytes used for increment the objectid.

Q27. Write syntax to create or select a database in MongoDB?

MongoDB uses DATABASE_NAME is used to create a database. The syntax of use DATABASE statement is as follows: **use DATABASE_NAME**

Q28. What is a collection in MongoDB?

In MongoDB, a **collection** is a grouping of documents such that documents within a collection can have different fields. A collection is the equivalent of a table in a relational database system and exists within a single database.

Q29. What is use of insertOne and insertMany in MongoDB?

In MongoDB, **insertOne()** is a method that is used whenever you want to insert a single document into your collection. It is basically used to insert a single document or record into the database. Whereas The **insertMany()** method inserts one or more documents in the collection. It takes an array of documents to insert into the collection.

Q30. What is sharding in MongoDB?

In MongoDB, **Sharding** is the process of distributing data across multiple hosts that is achieved by splitting large data sets into small data sets across multiple MongoDB instances. Basically, Sharding is used to store data across multiple machines.

Q31. What is writeConcern in MongoDB?

In MongoDB, **write concern** describes the level of acknowledgment requested from MongoDB for write operations to a standalone MongoDB or to replica sets or to sharded clusters, such that in sharded clusters, mongoose instances will pass the write concern on to the shards.

Q32. What is use of upsert in MongoDB?

MongoDB supports a feature known as **upsert** that mixes the functionality of an update with an insert. Upsert is an option that is used for update operation e.g. update(), findAndModify(), etc.

Q33. Explain what is Mongoose?

Mongoose is an Object Data Modeling or ODM library for MongoDB and Node. js. that manages relationships between data, provides schema validation and is used to translate between objects in code and the representation of those objects in MongoDB.

Q34. List some alternatives of MongoDB?

Some of the alternatives of MongoDB are:

- **JAM Stack -** Fast, secure, and dynamic websites served without web servers.
- **PostgreSQL -** SQL database is known for its reliability, features, and performance.
- **DynamoDB -** NoSQL database created by Amazon Web Services (AWS)

Q35. What is 32-bit nuances?

In computer **32-bit nuances** refer to the following:

- There is extra memory-mapped file activity with journaling.
- This will further constrain the limited db size of 32 bit builds.
- Thus, for now, journaling by default is disabled on 32 bit systems.

Q36. What is __v field in Mongoose?

In Mongoose the “_v” field is the **versionKey** is a property set on each document when first created by Mongoose. This is a document inserted through the mongo shell in a collection and this key-value contains the internal revision of the document.

Q37. What is Projection Operators? List some Projection Operators available in MongoDB?

In MongoDB, projection operators are used for selecting only the necessary data rather than selecting the whole of the data of a document. Some of the Projection Operators available in MongoDB are:

- **$ -** Projects the first element in an array that matches the query condition.
- **$elemMatch -** Projects the first element in an array that matches the specified $elemMatch condition.
- **$meta -** Projects the available per-document metadata.
- **$slice -** Limits the number of elements projected from an array.

Q38. What is Bitwise Query Operator in MongoDB?

In MongoDb, Bitwise Query Operators are used to match numeric or binary values in which any bit from a set of bit positions has a value of 0 and 1.

Q39. Explain the structure of ObjectID in MongoDB?

In MongoDB, the **structure of an ObjectId** is consists of a 12-byte BSON type. This 12-byte structure can be distinguished as the first 4 bytes of the ObjectId representing the time in seconds since the UNIX epoch. Then the next 3 bytes of the ObjectId represent the machine identifier which is then followed by the next 2 bytes of the ObjectId represent the process ID.

Q40. What are Indexes in MongoDB?

In MongoDB, **Indexes** are a special data structure that supports the efficient execution of queries. It is used to store a small portion of the collection's data set in an easy way in order to traverse it. Indexes improve the speed of search operations in the database because instead of searching the whole document, the search is performed on the indexes that hold only a few fields.

Q41. What is a covered query in MongoDB?

In MongoDB, a **covered query** is a query that can be satisfied entirely using an index and does not have to examine any documents. Covered queries help us in querying data faster by ensuring the index created contains all the fields required by the query and it doesn't require examining any documents apart from the indexed ones.

Q42. What is Aggregation in MongoDB?

In MongoDb, the **Aggregation operator** is used to group the values from multiple documents together, such that it can perform a variety of operations on the grouped data to return a single result for example sum, average, minimum, maximum, etc are some of the operations performed on the grouped data to return a computed result. MongoDB provides three ways to perform aggregation: the aggregation pipeline, the map-reduce function, and single-purpose aggregation methods.

Q43. What is oplog?

The **oplog** stands for operations log which is a special capped collection that keeps a rolling record of all operations that modify the data stored in your databases. Oplog is a log of every internal operation used for replication in a MongoDB cluster.

Q44. What is Replication in MongoDB?

In MongoDB, **replication** is the process of creating a copy of the same data set in more than one MongoDB server that can be achieved by using a Replica Set. Where a replica set is a group of mongodb processes that maintain the same data set.

Q45. What is a Storage Engine in MongoDB?

In MongoDB, the **storage engine** is the component of the database that is responsible for managing how data is stored, both in memory and on disk. It is based on memory-mapped files and manages BSON data in memory and on disk to support read and write operations.

Q46. Explain what is GridFS?

In MongoDb, **GridFS** is a specification for storing and retrieving files that exceed the BSON-document size limit of 16 MB. It is used to store and retrieve large files such as images, audio files, video files, etc.

Q47. What is use of capped collection in MongoDB?

Capped collections are fixed-size collections that restrict updates to the documents if the update results in increased document size and it ensures that the document size does not increase the size allocated on the disk.

Q48. What are Primary and Secondary Replica sets?

A **replica set** is a group of MongoDB instances that host the same data set. The primary is the only member in the replica set that receives write operations. Whereas Secondary members replicate this log and apply the operations to their data sets.

Q49. What is splitting in mongodb?

In MongoDb, **Splitting** is a process that keeps chunks from growing too large with the help of split command(). MongoDB splits the chunk, when a chunk grows beyond a specified chunk size, on the basis of shard key values the chunk represents.