

[Q33-Q55 Best Quality MongoDB C100DBA Exam Questions BraindumpsIT Realistic Practice Exams [2022]



Best Quality MongoDB C100DBA Exam Questions BraindumpsIT Realistic Practice Exams [2022]
Critical Information To MongoDB Certified DBA Associate Exam (Based on MongoDB 4.4) Pass the First Time

What is the registration procedure of MongoDB C100DBA Exam

Once the exam date has been announced by MongoDB, registration can be done. Registering for this exam is a little different from other certification exams. The steps for taking this exam must be followed. Registration for this exam will open on the date of the announcement. Registering for this certification exam is a free-of-cost process. This is a computer-based exam and is usually two hours long. **MongoDB C100DBA Dumps** is the best source for passing this certification exam with a high score. The exam takes place at Pearson Vue testing centers. This entry-level certification is only available in the English language. MongoDB website offers information about its training, certification, and more.

Q33. Which of the following is incorrect statement about find and findOne operations in MongoDB?

- * findQ returns all the documents in a collection while findOne() retrieves only the first one.
- * find.limit(1) is not the same query as findOne()
- * findOneQ returns the actual first document retrieved from a collection
- * findQ and findOneQ returns cursors to the collection documents

Q34. We can insert multiple documents in bulk using which of the following operations:

- * initializeUnorderedBulk
- * initializeUnorderedBulkOp
- * initializeBulk
- * initializeBulkOp

Q35. As per the aggregation pipeline optimization concepts, if you have a \$sort followed by a \$match:

- * Providing these parameters in any order does not impact the performance
- * \$sort moves before \$match
- * MongoDB does not do any movements by default and will use the order provided
- * \$match moves before \$sort

Q36. Which command is used to determine storage capacity of specific database?

- * mongotop
- * dbstats
- * mongostat
- * constats

Q37. In a collection that contains 100 post documents, what does the following command do? db.posts.

`find().skip(5).limit(5)`

- * Skip and limit nullify each other. Hence returning the first five documents
- * Limits the first five documents and then return them in reverse order
- * Skips the first five documents and returns the sixth document five times
- * Skips the first five documents and returns the next five

Q38. By default, the MongoDB cursor in mongo shell is configured to return how many documents? To get the next set of documents, which command is used?

- * 200, more
- * No limit, none
- * 20, it
- * 50, it

Q39. MongoDB is

- * None of the above
- * Object-oriented DBMS
- * Relational DBMS
- * Document-oriented DBMS

Q40. Update If Correct is an approach for which of the following concepts in MongoDB:

- * Concurrency Control
- * Atomicity
- * Performance Management
- * Transaction Management

Q41. Dada una coleccion, cuales devuelve con la siguiente query

`db.coleccion.find({'nombre':'ruben','apellido':'gomez'},{'nombre:1,apellido:1,aficion:1});`

- * { '_id': ObjectId('580a42b5dfb1b5a17427d302'), 'nombre': 'ruben' }

```
&#8220;ruben&#8221;, &#8220;apellido&#8221; : &#8220;gomez&#8221;, &#8220;aficion&#8221; :
```

```
v u &#8220;flipar&#8221; }
```

```
* { &#8220;_id&#8221; : ObjectId(&#8220;580a42acdfb5a17427d301&#8221;), &#8220;nombre&#8221; : &#8220;Luis&#8221;, &#8220;apellido&#8221; : &#8220;gomez&#8221;, &#8220;aficion&#8221; : u
```

```
&#8220;flipar&#8221; }
```

```
* { &#8220;_id&#8221; : ObjectId(&#8220;580a42acdfb5a17427d301&#8221;), &#8220;nombre&#8221; : &#8220;ruben&#8221;, &#8220;apellido&#8221; : &#8220;Pablo&#8221;, &#8220;aficion&#8221; : u
```

```
&#8220;flipar&#8221; }
```

```
* { &#8220;_id&#8221; : ObjectId(&#8220;580a42acdfb5a17427d301&#8221;), &#8220;nombre&#8221; : &#8220;ruben&#8221;, &#8220;apellido&#8221; : &#8220;gomez&#8221; >
```

Q42. Which is the correct order (lowest to highest) in which MongoDB compares the BSON types?

- * Number, Null, String and Object
- * Null, Number, Object and String
- * Null, Number, String and Object
- * String, Null, Number and Object

Q43. Consider the following example document from the sample collection. All documents in this collection have the same schema.

```
{ "_id" : 3,
  "a" : 7,
  "b" : 4
}
```

Which of the following queries will replace this with the document.

```
{
  "_id" : 7,
  "c" : 4,
  "b" : 4
}
```

- * `db.sample.update({ “_id” : 3 }, { “$set” : { “_id” : 7, “c” : 4 > > })`
- * `db.sample.update({ “_id” : 3 >, { “_id” : 7, “c” : 4 >, { “justOne” : true > } / This operation cannot be done with a single query.`
- * `db.sample.update({ “_id” : 3 >, { “_id” : 7, “c” : 4, { “$unset” : [“a”, “b”] } })`
- * `db.sample.update({ “_id” : 3 >, { “_id” : 7, “c” : 4 > })`

Q44. The oplog (operations log) is a special capped collection that keeps a rolling record of all operations that modify the data stored in your databases. All the replica set members contain a copy of the oplog in the following collection:

- * `<replicasetid>.oplog.rs`
- * `oplog.rs`
- * `<database>..oplog.rs`

* local.oplog.rs

Q45. Below is a sample document of `orders` collection

```
{
  cust_id: "abc123",
  ord_date: ISODate("2012-11-02T17:04:11.102Z"),
  status: 'A',
  price: 50,
  items: [ { sku: "xxx", qty: 25, price: 1 },
           { sku: "yyy", qty: 25, price: 1 } ]
}
```

Select operators for the below query to determine the sum of "qty" fields associated with the "cust_id".

```
db.orders.aggregate( [
  { $OPR1: "$items" },
  {
    $OPR2: {
      _id: "$cust_id",
      qty: { $OPR3: "$items.qty" }
    }
  }
] )
OPR2 is
```

\$sort

Q46. In a replica set, a _____ number of members ensures that the replica set is always able to select a primary.

- * Even
- * Odd
- * 2
- * Depends on the application architecture

Q47. Which of the following is true of the mechanics of replication in MongoDB? Check all that apply.

- * Members of a replica set may replicate data from any other data-bearing member of the set by default
- * Clients read from the nearest member of a replica set by default
- * Operations on the primary are recorded in a capped collection called the oplog

Q48. What does the `totalKeysExamined` field returned by the `explain` method indicate?

- * Number of documents that match the query condition
- * Number of index entries scanned
- * Details the completed execution of the winning plan as a tree of stages
- * Number of documents scanned

Q49. Consider the following document:

```
> db.c.find()
```

```
{ _id: 12, b: [ 3, 5, 7, 2, 1, -4, 3, 12 ] }
```

Which of the following queries on the `collection` will return only the first five elements of the array in the `b` field? E.g.,

Document you want returned by your query:

```
{ _id: 12, b: [ 3, 5, 7, 2, 1 ] }>  
* db.c.find( { b: { $slice: [ 0, 5 ] } } )  
* db.c.find( { b: [ 0, 5 ] } )>  
* db.c.find( { b: { $substr: [ 0, 5 ] } } )>>  
* db.c.find( { b: [ 0, 1, 2, 3, 4, 5 ] } )>  
* db.c.find( { b: [ 0, 5 ] } )>
```

Q50. Which of the following is a valid Replica Set configuration if you want:

- 1- Have 3 copies of everything
- 2- That RS3 is never primary
- 2- That RS1 and RS2 can be primary?

You had to see the different configurations, RS3 could be hidden or priority 0 (But not a referee because we need

3 replicas), while RS1 and RS2 could NOT have priority 0 or be hidden or anything like that In a 4-member RS RSO, RS1, RS2 and RS3 + Referee, RSO (primary) falls after some write operations that have replicated RS1 and RS2 (but NOT RS3), who can get up as the new primary?

The configuration comes and in it we see that RS2 has a `hidden: true` (or a `priority: 0`, (I don't remember)

- * ORS1
- * ORS2
- * ORS3
- * O arbiter
- * RSO

Q51. Consider the following posts document:

```
{  
  _id: 1,  
  post_text: 'This is my first post',  
  author: 'Tom',  
  tags: [ 'tutorial', 'quiz', 'facebook', 'learning', 'fun' ]  
}
```

Which of the following queries will return the documents but with only the first two tags in the tags array?

```
* db.posts.find( { author: 'Tom' }.limit( { tags: 2 } )
```

- * `db.posts.find({author:'Tom'}).limit($slice:{tags:2})`
- * Both `db.posts.find({author:'Tom'}, {tags:{$slice:2}})` and

`db.posts.find({author:'Tom'}).limit($slice: {tags:2})` are valid. `$slice` works both with projection and limit.

- * `db.posts.find({author:'Tom'},>,{tags:{$slice:2}})`

Q52. To add a new user and enable authentication in MongoDB, which of the following steps need be executed?

- * update users collection and restart mongod
- * All of the above
- * update users collection and restart mongod with `-auth` option
- * update users collection and run `db.enableAuthenticationQ`

Q53. Given a collection posts as shown below having a document array comments, which of the following command will create an index on the comment author descending?

```
{
  _id:1,
  post_text:This is a sample post,
  author:Tom,
  comments:[
    {
      author:Joe,
      comment_text:This is comment 1
    },
    {
      author:Leo,
      comment_text:This is comment 2
    }
  ]
}
```

- * `db.posts.createIndex({comment.author:-1});`
- * `db.posts.createIndex({comments.$comment.author:-1});`
- * `db.posts.createIndex({comments.author:-1});`

Q54. Which of the following needs to be performed prior to initiate backup on a sharded cluster?

- * `db.stopBalancer()`
- * `sh.stopServer()`
- * `db.stopServer()`
- * `sh.stopBalancer()`

Q55. What read preference should your application use if you want to read from the primary under normal circumstances but allow reads from secondaries when a primary is unavailable?

- * `secondaryPreferred`
- * `Secondary`
- * `nearest`
- * `primary`
- * `primaryPreferred`

For more info read reference
MongoDB C100DBA Exam

Certification Path of MongoDB C100DBA Exam - MongoDB Developer Certified Associate- MongoDB Certified Developer- MongoDB Certified Engineer- MongoDB Certified Architect **C100DBA EXAM DUMPS WITH GUARANTEED SUCCESS:** https://www.braindumpsit.com/C100DBA_real-exam.html]