

Scheme of Evaluation
Internal Assessment Test 3 – May 2024

Sub:	NoSQL Database					Code:	18CS823		
Date:	11/5/2024	Duration:	90 mins	Max Marks:	50	Sem:	VIII	Branch:	ISE

Note: Answer Any five full questions.

Question #		Description	Marks Distribution		Max Marks
1	a)	What is document database? Give examples of any document that can be stored into it and explain. Document database Example and explanation.	2M 3M	5M	5M
1	b)	Elaborate the suitable use cases of document databases. Explanation	5M	5M	5M
2		With examples differentiate between queries written in SQL and its equivalent query in Mongo Shell. Definition Explanation	4M 6M	10M	10M
3		With suitable diagrams, discuss and analyze-- a) Horizontal sharding in MongoDB for adding a new node to an existing replica set. b) Each shard is replica set. Diagram Analysis and discussion	3M 7M	10M	10M

4		<p>What are graph databases? Briefly describe the relationships in graph databases, with neat diagram</p> <p>Defination</p> <p>Explanation with Diagram</p>	<p>3M</p> <p>7M</p>	10M	10M
5		<p>With neat diagram, explain the three ways in which graph databases can be scaled.</p> <p>Diagram</p> <p>Explanation</p>	<p>3M</p> <p>7M</p>	10M	
6	a)	<p>Explain query features of graph database with examples</p> <p>Explanatin</p>	<p>5M</p>	5M	10M
6	b)	<p>Explain Consistency, Transaction and Availability with respect to graph database..</p> <p>Explanation</p> <p>Example</p>	<p>3M</p> <p>2M</p>	5M	

Scheme Of Evaluation
Internal Assessment Test 3 –May 2024

Sub:	NoSQL Database						Code:	18CS823	
Date:	11/5/2024	Duration:	90mins	Max Marks:	50	Sem:	VIII	Branch:	ISE

Note: Answer Any full five questions

Q-1 a) What is document database? Give examples of any document that can be stored into it and explain.

b) Elaborate the suitable use cases of document databases.

Ans-1

Documents are the main concept in document databases. The database stores and retrieves documents, which can be XML, JSON, BSON, and so on. These documents are self-describing, hierarchical tree data structures which can consist of maps, collections, and scalar values. The documents stored are similar to each other but do not have to be exactly the same. Document databases store documents in the value part of the key-value store; think about document

databases as key-value stores where the value is examinable. Let's look at how terminology compares in Oracle and MongoDB.

The `_id` is a special field that is found on all documents in Mongo, just like ROWID in Oracle. In MongoDB, `_id` can be assigned by the user, as long as it is unique.

1.1 What Is a Document Database?

```
{ "firstname": "Martin", "likes": [ "Biking", "Photography" ], "lastcity": "Boston", "lastVisited":  
}
```

The above document can be considered a row in a traditional RDBMS. Let's look at another document:

```
{ "firstname": "Prمود",  
"citiesvisited": [ "Chicago", "London", "Pune", "Bangalore" ], "addresses": [  
{ "state": "AK", "city": "DILLINGHAM",  
"type": "R" }, { "state": "MH", "city": "PUNE", "type": "R" } ], "lastcity": "Chicago" }
```

Looking at the documents, we can see that they are similar, but have differences in attribute names.

This is allowed in document databases. The schema of the data can differ across documents, but these documents can still belong to the same collection—unlike an RDBMS where every row in a table has to follow the same

schema. We represent a list of cities visited as an array, or a list of addresses as list of documents embedded inside the main document. Embedding child documents as sub objects inside documents provides for easy access and better performance.

If you look at the documents, you will see that some of the attributes are similar, such as firstname or city. At the same time, there are attributes in the second document which do not exist in the first document, such as addresses, while likes is in the first document but not the second.

This different representation of data is not the same as in RDBMS where every column has to be defined, and if it does not have data it is marked as empty or set to null. In documents, there are no empty attributes; if a given attribute is not found, we assume that it was not set or not relevant to the document. Documents allow for new attributes to be created without the need to define them or to change the existing documents.

- a) Some of the popular document databases we have seen are MongoDB [MongoDB], CouchDB [CouchDB], Terrastore [Terrastore], OrientDB [OrientDB], RavenDB [RavenDB], and of course the well-known and often reviled Lotus Notes [Notes Storage Facility] that uses document storage.

1.2 Suitable Use Cases

Event Logging

Applications have different event logging needs; within the enterprise, there are many different applications that want to log events. Document databases can store all these different types of events and can act as a central data store for event storage. This is especially true when the type of data being captured by the events keeps changing. Events can be sharded by the name of the application where the event originated or by the type of event such as order_processed or customer_logged.

Content Management Systems, Blogging Platforms

Since document databases have no predefined schemas and usually understand JSON documents, they work well in content management systems or applications for publishing websites, managing user comments, user registrations, profiles, web-facing documents.

Web Analytics or Real-Time Analytics

Document databases can store data for real-time analytics; since parts of the document can be updated, it's very easy to store page views or unique visitors, and new metrics can be easily added without schema changes.

E-Commerce Applications

E-commerce applications often need to have flexible schema for products and orders, as well as the ability to evolve their data models without expensive database refactoring or data migration (“[Schema Changes in a NoSQL Data Store](#),” p. 128).

When Not to Use

- a) There are problem spaces where document databases are not the best solution

Q-2 With examples differentiate between queries written in SQL and its equivalent query in Mongo Shell.

Ans- Queries in SQL and MongoDB's shell (using the Mongo query language) operate on different database systems with distinct data models and query syntax. Here's how queries in SQL and MongoDB's shell differ, along with examples:

In SQL, queries are structured using a structured query language designed for relational databases, such as MySQL, PostgreSQL, or SQLite. SQL queries typically involve tables with rows and columns, and they follow a standardized syntax. Here's an example SQL query:

```
SELECT name, age FROM users WHERE age > 25;
```

Explanation: This SQL query selects the name and age columns from the users table where the age is greater than 25.

In MongoDB's shell, queries are constructed using MongoDB's query language, which is specifically designed for document-oriented databases like MongoDB. MongoDB stores data in flexible, JSON-like documents, and queries are expressed as JSON-like structures. Here's the equivalent query in MongoDB's shell:

```
db.users.find({ age: { $gt: 25 } }, { name: 1, age: 1, _id: 0 });
```

Explanation: This MongoDB shell command queries the users collection for documents where the age field is greater than 25. It projects ({ name: 1, age: 1, _id: 0 }) only the name and age fields while excluding the _id field from the result.

Key differences between the two:

Syntax: SQL queries use a structured language with keywords like SELECT, FROM, and WHERE, while MongoDB's shell queries are expressed as JSON-like objects with field-value pairs and operators.

Data Model: SQL operates on tables with rows and columns, while MongoDB operates on collections of JSON-like documents.

Schema Flexibility: SQL databases typically have a fixed schema, while MongoDB allows for flexible schemas within collections.

Query Operators: MongoDB's query language provides a rich set of operators for querying nested documents and arrays, which may differ from the operators used in SQL.

In summary, while both SQL and MongoDB's shell queries achieve similar goals of retrieving data from a database, they differ in syntax, data model, and query operators due to their underlying database systems' design and principles.

Q-3 With suitable diagrams, discuss and analyze--

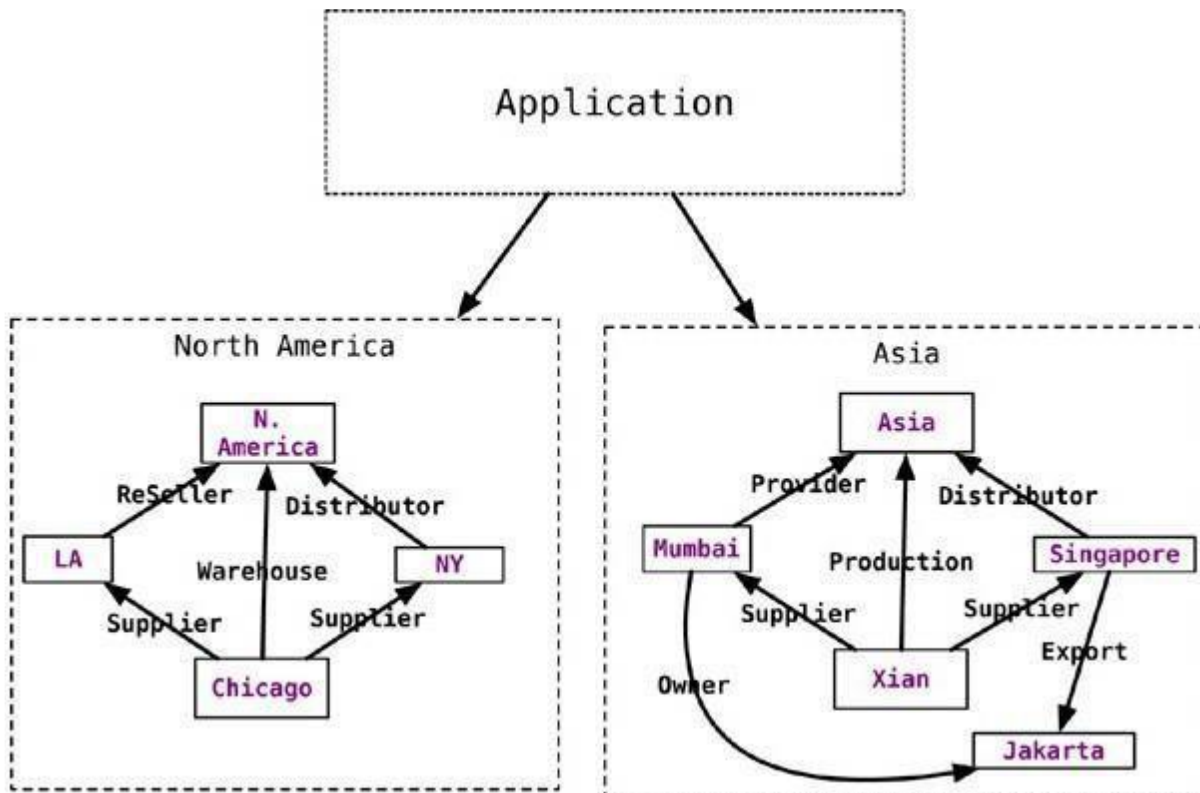
- a) Horizontal sharding in MongoDB for adding a new node to an existing replica set.**
- b) Each shard is replica set.**

In NoSQL databases, one of the commonly used scaling techniques is sharding, where data is split and distributed across different servers. With graph databases, sharding is difficult, as graph databases are not aggregate-oriented but relationship-oriented. Since any given node can be related to any other node, storing related nodes on the same server is better for graph traversal. Traversing a graph when the nodes are on different machines is not good for performance. Knowing this limitation of the graph databases, we can still scale them using some common techniques described by Jim Webber [Webber Neo4J Scaling].

Generally speaking, there are three ways to scale graph databases. Since machines now can come with lots of RAM, we can add enough RAM to the server so that the working set of nodes and relationships is held entirely in memory. This technique is only helpful if the dataset that we are working with will fit in a realistic amount of RAM.

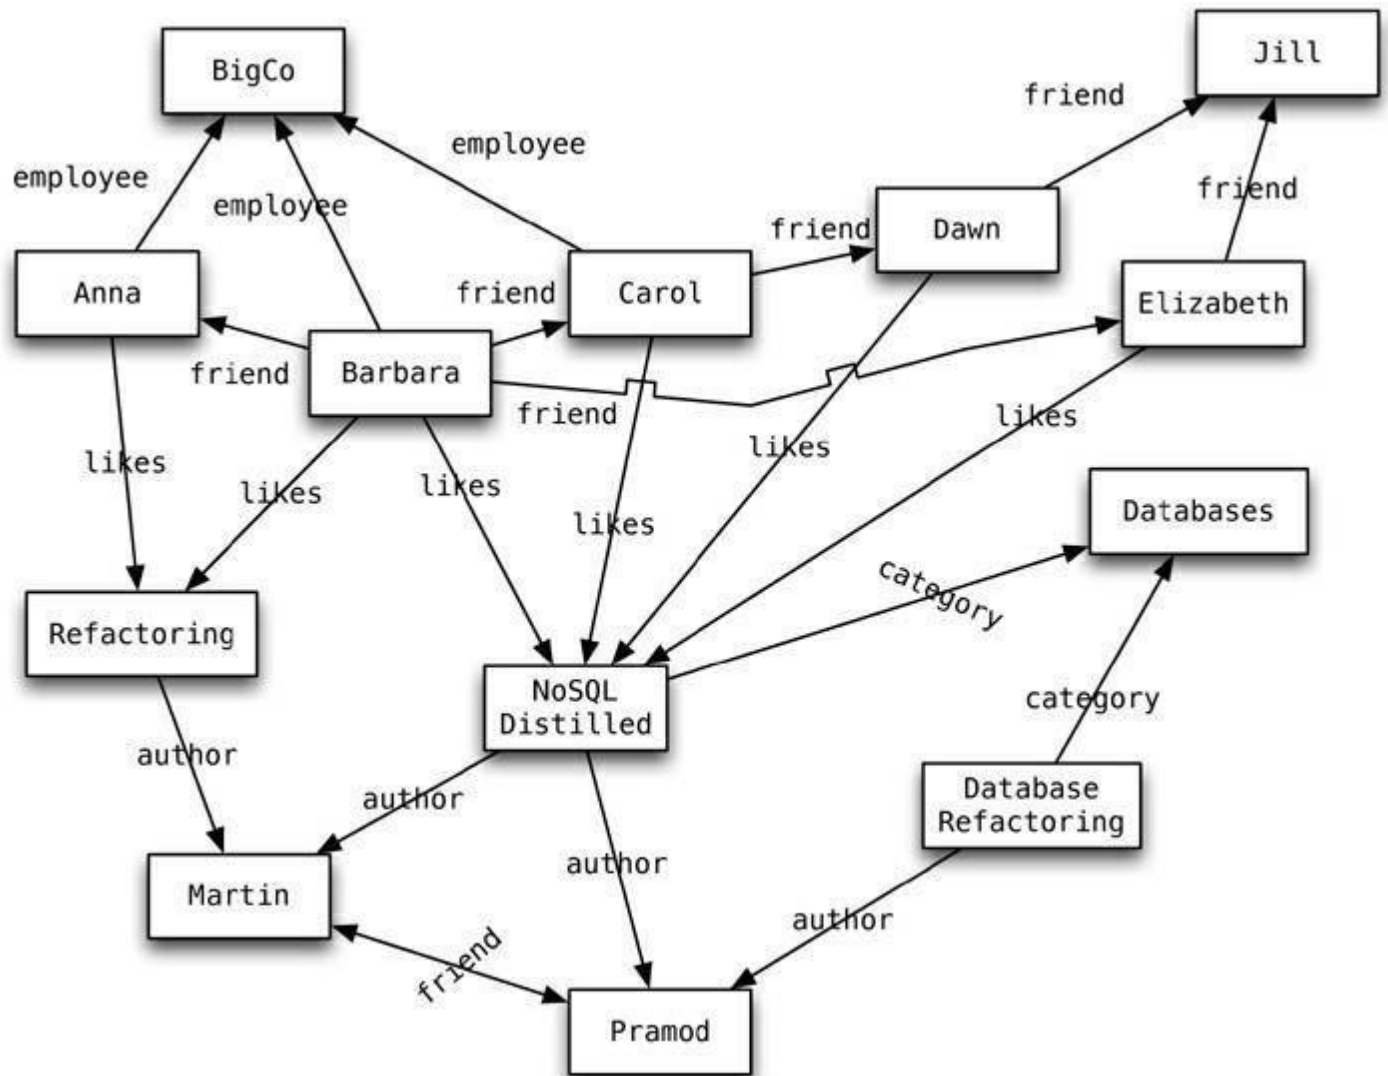
We can improve the read scaling of the database by adding more slaves with read-only access to the data, with all the writes going to the master. This pattern of writing once and reading from many servers is a proven technique in MySQL clusters and is really useful when the dataset is large enough to not fit in a single machine's RAM, but small enough to be replicated across multiple machines.

Slaves can also contribute to availability and read-scaling, as they can be configured to never become a master, remaining always read-only.



When the dataset size makes replication impractical, we can shard the data from the application side using domain-specific knowledge. For example, nodes that relate to the North America can be created on one server while the nodes that relate to Asia on another. This application-level sharding needs to understand that nodes are stored on physically different databases (Figure 5.3).

Q-4 What are graph databases? Briefly describe the relationships in graph databases, with neat diagram.




In the example graph in Figure 5.1, we see a bunch of nodes related to each other. Nodes are entities that have properties, such as name. The node of Martin is actually a node that has property of name set to Martin. We also see that edges have types, such as likes, author, and so on. These properties let us organize the nodes; for example, the nodes Martin and Pramod have an edge connecting them with a relationship type of friend. Edges can have multiple properties. We can assign a property of since on the friend relationship type between Martin and Pramod. Relationship types have directional significance; the friend relationship type is bidirectional but likes is not. When Dawn likes NoSQL Distilled, it does not automatically mean NoSQL Distilled likes Dawn. Once we have a graph of these nodes and edges created, we can query the graph in many ways, such as “get all nodes employed by Big Co that like NoSQL Distilled.” A query on the graph is also known as traversing the graph. An advantage of the graph databases is that we can change the traversing requirements without having to change the nodes or edges. If we want to “get all nodes that like NoSQL Distilled,” we can do so without having to change the existing data or the model of the database, because we can traverse the graph any way we like.

Usually, when we store a graph-like structure in RDBMS, it's for a single type of relationship ("who is my manager" is a common example). Adding another relationship to the mix usually means a lot of schema changes and data movement, which is not the case when we are using graph databases. Similarly, in relational databases we model the graph beforehand based on the Traversal we want; if the Traversal changes, the data will have to change.

In graph databases, traversing the joins or relationships is very fast. The relationship between nodes is not calculated at query time but is actually persisted as a relationship. Traversing persisted relationships is faster than calculating them for every query.

Nodes can have different types of relationships between them, allowing you to both represent relationships between the domain entities and to have secondary relationships for things like category, path, time-trees, quad-trees for spatial indexing, or linked lists for sorted access. Since there is no limit to the number and kind of relationships a node can have, all they can be represented in the same graph database.

Q-5 With neat diagram, explain the three ways in which graph databases can be scaled.



Ans- The idea of scaling is to add nodes or change data storage without simply migrating the database to a bigger box. We are not talking about making application changes to handle more load; instead, we are interested in what features are in the database so that it can handle more load.

Scaling for heavy-read loads can be achieved by adding more read slaves, so that all the reads can be directed to the slaves. Given a heavy-read application, with our 3-node replica- set cluster, we can add more read capacity to the cluster as the read load increases just by

adding more slave nodes to the replica set to execute reads with the `slaveOk` flag (Figure 1.2). This is horizontal scaling for reads.

Figure 1.2. Adding a new node, mongo D, to an existing replica-set cluster

Once the new node, mongo D, is started, it needs to be added to the replica set.`rs.add("mongod:27017");`

When a new node is added, it will sync up with the existing nodes, join the replica set as secondary node, and start serving read requests. An advantage of this setup is that we do not have to restart any other nodes, and there is no downtime for the application either.

When we want to scale for write, we can start sharding (“[Sharding](#),”) the data. Sharding is similar to partitions in RDBMS where we split data by value in a certain column, such as state or year. With RDBMS, partitions are usually on the same node, so the client application does not have to query a specific partition but can keep querying the base table; the RDBMS takes care of finding the right partition for the query and returns the data. In sharding, the data is also split by certain field, but then moved to different Mongo nodes. The data is dynamically moved between nodes to ensure that shards are always balanced. We can add more nodes to the cluster and increase the number of writable nodes, enabling horizontal scaling for writes.

```
db.runCommand( { shardcollection : "ecommerce.customer",key : { firstname : 1 } } )
```

Splitting the data on the first name of the customer ensures that the data is balanced across the shards for optimal write performance; furthermore, each shard can be a replica set ensuring better read performance within the shard (Figure 1.3). When we add a new shard to this existing sharded cluster, the data will now be balanced across four shards instead of three. As all this data movement and infrastructure refactoring is happening, the application will not experience any downtime, although the cluster may not perform optimally when large amounts of data are being moved to rebalance the shards.

Figure 1.3. MongoDB sharded setup where each shard is a replica set

The shard key plays an important role. You may want to place your MongoDB database shards closer to their users, so sharding based on user location may be a good idea. When

sharding by customer location, all user data for the East Coast of the USA is in the shards that are served from the East Coast, and all user data for the West Coast is in the shards that are on the West Coast.

Q-6 a) Explain query features of graph databases in with examples.

Ans- Document databases provide different query features. CouchDB allows you to query via views— complex queries on documents which can be either materialized (“[Materialized Views](#),” p. 30) or dynamic (think of them as RDBMS views which are either materialized or not). With CouchDB, if you need to aggregate the number of reviews for a product as well as the average rating, you could add a view implemented via map-reduce (“[Basic Map-Reduce](#),” p. 68) to return the count of reviews and the average of their ratings.

When there are many requests, you don't want to compute the count and average for every request; instead you can add a materialized view that precomputes the values and stores the results in the database. These materialized views are updated when queried, if any data was changed since the last update.

One of the good features of document databases, as compared to key-value stores, is that we can query the data inside the document without having to retrieve the whole document by its key and then introspect the document.

This feature brings these databases closer to the RDBMS query model.

MongoDB has a query language which is expressed via JSON and has constructs such as

`$query` for the where clause, `$orderby` for sorting the data, or `$explain` to show the execution plan of the query.

There are many more constructs like these that can be combined to create a MongoDB query.

Let's look at certain queries that we can do against MongoDB. Suppose we want to return all the documents in an order collection (all rows in the order table). The SQL for this would be:

```
SELECT * FROM order
```

The equivalent query in Mongo shell would be: `db.order.find()`

Selecting the orders for a single customerId of 883c2c5b4e5b would be: `SELECT * FROM order WHERE customerId = "883c2c5b4e5b"`

The equivalent query in Mongo to get all orders for a single customerId of 883c2c5b4e5b:

```
db.order.find({"customerId":"883c2c5b4e5b"})
```

Similarly, selecting orderId and orderDate for one customer in SQL would be: `SELECT orderId,orderDate FROM order WHERE customerId = "883c2c5b4e5b"`

and the equivalent in Mongo would be: `db.order.find({customerId:"883c2c5b4e5b"},{orderId:1,orderDate:1})`

Similarly, queries to count, sum, and so on are all available. Since the documents are aggregated objects, it is really easy to query for documents that have to be matched using the fields

with child objects. Let's say we want to query for all the orders where one of the items ordered has a name like Refactoring. The SQL for this requirement would be:

```
SELECT * FROM customerOrder, orderItem, product WHERE customerOrder.orderId =  
orderItem.customerOrderId AND orderItem.productId = product.productId AND product.name LIKE  
'%Refactoring%'
```

and the equivalent Mongo query would be: `db.orders.find({"items.product.name":/Refactoring/})`

The query for MongoDB is simpler because the objects are embedded inside a single document and you can query based on the embedded child documents.

b) Explain Consistency, Transaction and Availability with respect to graph database.

Ans- b) While there are many specialized document databases, we will use MongoDB as a representative of the feature set. Keep in mind that each product has some features that may not be found in other document databases. Let's take some time to understand how MongoDB works. Each MongoDB instance has multiple *databases*, and each database can have multiple *collections*. When we compare this with RDBMS, an RDBMS instance is the same as MongoDB instance, the schemas in RDBMS are similar to MongoDB databases, and the RDBMS tables are collections in MongoDB. When we store a document, we have to choose which database and collection this document belongs in for example, `database.collection.insert(document)`, which is usually represented as `db.coll.insert(document)`.

Consistency

Consistency in MongoDB database is configured by using the replica sets and choosing to wait for the writes to be replicated to all the slaves or a given number of slaves. Every write can specify the number of servers the write has to be propagated to before it returns as successful.

A command like `db.runCommand({ getlasterror : 1 , w : "majority" })` tells the database how strong is the consistency you want. For example, if you have one server and specify the `w` as `majority`, the write will return immediately since there is only one node. If you have three nodes in the replica set and specify `w` as `majority`, the write will have to complete at a minimum of two nodes before it is reported as a success. You can increase the `w` value for stronger consistency but you will suffer on write performance, since now the writes have to complete at more nodes.

Replica sets also allow you to increase the read performance by allowing reading from slaves by setting `slaveOk`; this parameter can be set on the connection, or database, or collection, or individually for each operation.

```
Mongo mongo = new Mongo("localhost:27017"); mongo.slaveOk();
```

Here we are setting `slaveOk` per operation, so that we can decide which operations can work with data from the slave node.

```
DBCcollection collection = getOrderCollection(); BasicDBObject query = new BasicDBObject();
query.put("name", "Martin");
DBCcursor cursor = collection.find(query).slaveOk();
```

Similar to various options available for read, you can change the settings to achieve strong write consistency, if desired. By default, a write is reported successful once the database receives it; you can change this so as to wait for the writes to be synced to disk or to propagate to two or more slaves. This is known as WriteConcern: You make sure that certain writes are written to the master and some slaves by setting WriteConcern to

REPLICAS_SAFE. Shown below is code where we are setting the WriteConcern for all writes to a collection:
DBCcollection shopping = database.getCollection("shopping"); shopping.setWriteConcern(REPLICAS_SAFE);

Write Concern can also be set per operation by specifying it on the save command:

```
WriteResult result = shopping.insert(order, REPLICAS_SAFE);
```

There is a tradeoff that you need to carefully think about, based on your application needs and business requirements, to decide what settings make sense for slaveOk during read or what safety level you desire during write with WriteConcern.

Transactions

Transactions, in the traditional RDBMS sense, mean that you can start modifying the database with insert, update, or delete commands over different tables and then decide if you want to keep the changes or not by using commitor rollback. These constructs are generally not available in NoSQL solutions—a write either succeeds or fails. Transactions at the single-document level are known as **atomic transactions**. Transactions involving more than one operation are not possible, although there are products such as RavenDB that do support transactions across multiple operations.

By default, all writes are reported as successful. A finer control over the write can be achieved by using WriteConcern parameter. We ensure that order is written to more than one node before it's reported successful by using WriteConcern.REPLICAS_SAFE. Different levels of WriteConcern let you choose the safety level during writes; for example, when writing log entries, you can use lowest level of safety, WriteConcern.NONE.

```
final Mongo mongo = new Mongo(mongoURI); mongo.setWriteConcern(REPLICAS_SAFE);
DBCcollection shopping = mongo.getDB(orderDatabase).getCollection(shoppingCollection); try { WriteResult
result = shopping.insert(order, REPLICAS_SAFE);
//Writes made it to primary and at least one secondary
} catch (MongoException writeException) {
//Writes did not make it to minimum of two nodes including primary dealWithWriteFailure(order,
writeException);
}
```

Availability

The CAP theorem (“The CAP Theorem,” p. 53) dictates that we can have only two of Consistency, Availability, and Partition Tolerance. Document databases try to improve on availability by replicating data using the master-slave setup. The same data is available on multiple nodes and the clients can get to the data even when the primary node is down. Usually, the application code does not have to determine if the primary node is available or not. MongoDB implements replication, providing high availability using replica sets.

In a replica set, there are two or more nodes participating in an asynchronous master-slave replication. The replica-set nodes elect the master, or primary, among themselves. Assuming all the nodes have equal voting rights, some nodes can be favored for being closer to the other servers, for having more RAM, and so on; users can affect this by assigning a priority a number between 0 and 1000 to a node.

All requests go to the master node, and the data is replicated to the slave nodes. If the master node goes down, the remaining nodes in the replica set vote among themselves to elect a new master; all future requests are routed to the new master, and the slave nodes start getting data from the new master. When the node that failed comes back online, it joins in as a slave and catches up with the rest of the nodes by pulling all the data it needs to getcurrent.

[Figure 1.1](#) is an example configuration of replica sets. We have two nodes, **mongo A** and **mongo B**, running the MongoDB database in the primary data-center, and **mongo C** in the secondary datacenter. If we want nodes in the primary datacenter to be elected as primary nodes, we can assign them a higher priority than the other nodes. More nodes can be added to the replica sets without having to take them offline.