

UNIT II NOSQL DATA MANAGEMENT

Introduction to NoSQL – aggregate data models – key-value and document data models – relationships – graph databases – schemaless databases – materialized views – distribution models – master-slave replication – consistency - Cassandra – Cassandra data model – Cassandra examples – Cassandra clients

2.1 Introduction to NoSQL

NoSQL, known as Not only SQL database, provides a mechanism for storage and retrieval of data and is the next generation database . It has a distributed architecture with MongoDB and is open source. Most of the NoSQL are open source and it has a capability of horizontal scalability which means that commodity kind of machines could be added

The capacity of your clusters can be increased. It is schema free and there is no requirement to design the tables and pushing the data to it. NoSQL provides easy replication claiming there are very less manual interventions in this. Once the replication is done, the system will automatically take care of fail overs.

The crucial factor about NoSQL is that it can handle huge amount of data and can achieve performance by adding more machines to your clusters and can be implemented on commodity hardware. There are close to 150 NoSQL databases in the market which will make it difficult to choose to choose the right pick for your system.

Why NoSQL?

There are different kinds of data that are structured, unstructured and semi-structured and hence RDBMS systems are not designed to manage these types of data in an effecient way. NoSql databases comes in to the picture and are capable to manage it . People today are use different kinds of methodology and if you talk about the code velocity and implementation level, people wish to execute various activities at the same time. Talking about the traditional way, it was quite difficult in terms of taking care of the database and writing your application according to the database. So with the new NoSQL, things have become easier.

Benefits of NoSQL

NoSQL claims various benefits associated with its usage. From managing data to scalability, it use has increased drastically. Huge amount of data can be managed and then streamed. Different kinds of data like structured and semi structured data can be analyzed and analytical activities can be performed.

If you're not able to drive value out of any data, you can capture this kind of information. It's object oriented programming is easy to use and flexible. It provides horizontal scaling instead of expensive hardware. If you want to scale vertically, then additional CPU's and RAMs have to be purchased. But here horizontal scaling can be used with the help of commodity hardware.

NoSQL Database Categories

Document Database– It pairs each key with a complex data structure known as document. It can contain many different key value pairs, or key array pairs or even nested documents

Key value stores– They are the simplest NoSQL databases. Every single item in the database is stored as an attribute name or key together with its value.

Graph store– They are used to store information about networks, such as social connections. Graph stores include Neo4J and HyperGraphDB.

Wide column stores- Wide column stores such as Cassandra and HBase are optimized for queries over large datasets, and store columns of data together, instead of rows.

NoSQL Vs SQL Comparison

ACID Property

ACID is a set of properties that guarantees that are processed reliably. In the context of databases, a single logical operation on the data is called a transaction.

Atomicity- Atomicity defines transaction as a logical unit of work which must be either completed with all of its data modifications or none of them is performed.

Consistency- At the very end of the transaction, all data must be left in a consistent state.

Isolation- Modifications of data performed by a transaction must be independent of another transaction. Unless this happens, the outcome of a transaction may be erroneous.

Durability- When the transaction is completed, effects of the modifications performed by the transaction must be permanent in the system.

2.2 Aggregate data models

Aggregate means a collection of objects that are treated as a unit. In NoSQL Databases, an aggregate is a collection of data that interact as a unit. Moreover, these units of data or aggregates of data form the boundaries for the ACID operations.

Aggregate Data Models in NoSQL make it easier for the Databases to manage data storage over the clusters as the aggregate data or unit can now reside on any of the machines. Whenever data is retrieved from the Database all the data comes along with the Aggregate Data Models in NoSQL.

Aggregate Data Models in NoSQL don't support ACID transactions and sacrifice one of the ACID properties. With the help of Aggregate Data Models in NoSQL, you can easily perform OLAP operations on the Database.

You can achieve high efficiency of the Aggregate Data Models in the NoSQL Database if the data transactions and interactions take place within the same aggregate.

Types of Aggregate Data Models in NoSQL Databases

The Aggregate Data Models in NoSQL are majorly classified into 4 Data Models listed below:

Key-Value Model

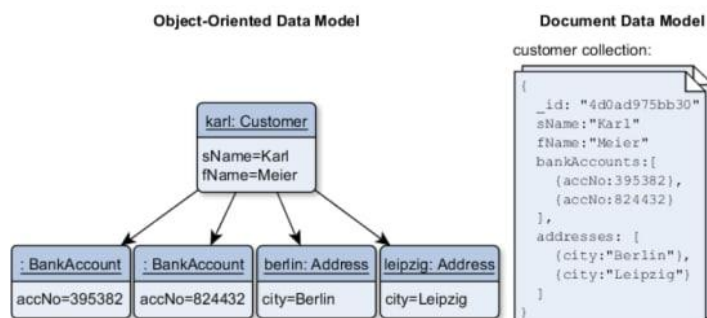
| Key | Value |
|-----|------------------|
| K1 | AAA,BBB,CCC |
| K2 | AAA,BBB |
| K3 | AAA,DDD |
| K4 | AAA,2,01/01/2015 |
| K5 | 3,ZZZ,5623 |

The Key-Value Data Model contains the key or an ID used to access or fetch the data of the aggregates corresponding to the key. In this Aggregate Data Models in NoSQL, the data of the aggregates are secure and encrypted and can be decrypted with a Key.

Use Cases:

- These Aggregate Data Models in NoSQL Database are used for storing the user session data.
- Key Value-based Data Models are used for maintaining schema-less user profiles.
- It is used for storing user preferences and shopping cart data.

Document Model

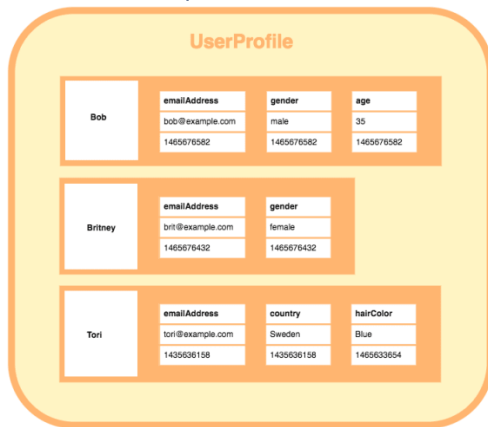


The Document Data Model allows access to the parts of aggregates. In this Aggregate Data Models in NoSQL, the data can be accessed in an inflexible way. The Database stores and retrieves documents, which can be XML, JSON, BSON, etc. There are some restrictions on data structure and data types of the data aggregates that are to be used in this Aggregate Data Models in NoSQL Database.

Use Cases:

- Document Data Models are widely used in E-Commerce platforms
- It is used for storing data from content management systems.
- Document Data Models are well suited for Blogging and Analytics platforms.

Column Family Model



Column family is an Aggregate Data Models in NoSQL Database usually with big-table style Data Models that are referred to as column stores. It is also called a two-level map as it offers a two-level aggregate structure. In this Aggregate Data Models in NoSQL, the first level of the Column family contains the keys that act as a row identifier that is used to select the aggregate data. Whereas the second level values are referred to as columns.

Use Cases:

- Column Family Data Models are used in systems that maintain counters.
- These Aggregate Data Models in NoSQL are used for services that have expiring usage.
- It is used in systems that have heavy write requests.

Graph-Based Model



Graph-based data models store data in nodes that are connected by edges. These Aggregate Data Models in NoSQL are widely used for storing the huge volumes of complex aggregates and multidimensional data having many interconnections between them.

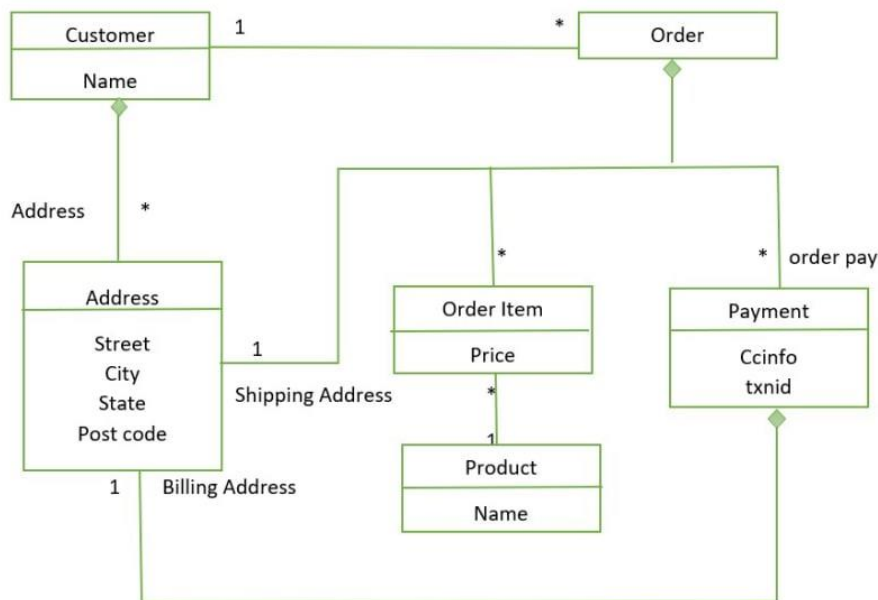
Use Cases:

- Graph-based Data Models are used in social networking sites to store interconnections.
- It is used in fraud detection systems.
- This Data Model is also widely used in Networks and IT operations.

Steps to Build Aggregate Data Models in NoSQL Databases

Now that you have a brief knowledge of Aggregate Data Models in NoSQL Database. In this section, you will go through an example to understand how to design Aggregate Data Models in NoSQL. For this, a Data Model of an E-Commerce website will be used to explain Aggregate Data Models in NoSQL.

This example of the E-Commerce Data Model has two main aggregates – customer and order. The customer contains data related to billing addresses while the order aggregate consists of ordered items, shipping addresses, and payments. The payment also contains the billing address.



If you notice a single logical address record appears 3 times in the data, but its value is copied each time wherever used. The whole address can be copied into an aggregate as needed. There is no pre-defined format to draw the aggregate boundaries. It solely depends on whether you want to manipulate the data as per your requirements.

The Data Model for customer and order would look like this.

```

// in customers
{
  "customer": {
    "id": 1,
    "name": "Martin",
    "billingAddress": [{"city": "Chicago"}],
    "orders": [
      {
        "id": 99,
        "customerId": 1,
        "orderItems": [
          {
            "productId": 27,
            "price": 32.45,
            "productName": "NoSQL Distilled"
          }
        ],
        "shippingAddress": [{"city": "Chicago"}]

        "orderPayment": [
          {
            "ccinfo": "1000-1000-1000-1000",
            "txnId": "abelif879rft",
            "billingAddress": {"city": "Chicago"}
          }
        ]
      }
    ]
  }
}

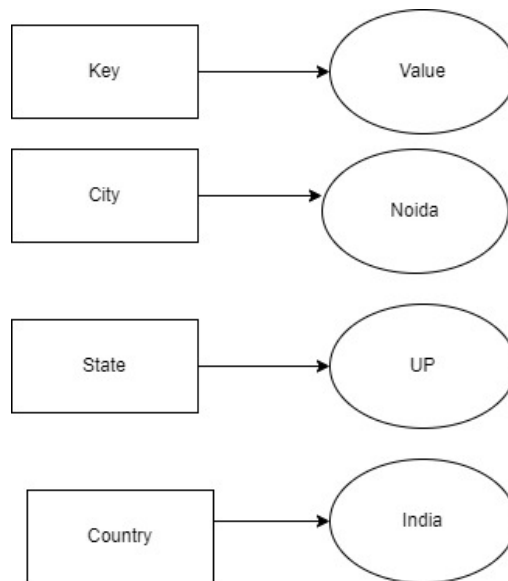
```

In these Aggregate Data Models in NoSQL, if you want to access a customer along with all customer's orders at once. Then designing a single aggregate is preferable. But if you want to access a single order at a time, then you should have separate aggregates for each order. It is very content-specific

2.2.1 Key-value and document data models

Key-Value Data Model:

The key-value data model is one of the simplest and most fundamental data models used in NoSQL databases. In this model, data is stored as a collection of key-value pairs, where each key is a unique identifier and each value is associated with that key. The database retrieves and stores data based on its corresponding key.



Key-value databases provide a highly efficient way to store and retrieve data, as accessing a specific value is fast since it is directly linked to its key. These databases are often optimized for high-speed reads and writes, making them suitable for caching, session management, and simple data storage needs.

Some characteristics of key-value data models are as follows:

1. **Simplicity:** Key-value databases offer a straightforward and simple data model, making them easy to understand and use.
2. **Scalability:** Key-value databases are designed for high scalability by allowing data to be distributed across multiple servers or clusters.
3. **Flexibility:** The values in key-value pairs can vary in structure, meaning different keys can have different data types and sizes.

Popular key-value databases include Redis, Riak, and Amazon DynamoDB.

How do key-value databases work?

A number of easy strings or even a complicated entity are referred to as a value that is associated with a key by a key-value database, which is utilized to monitor the entity. Like in many programming paradigms, a key-value database resembles a map object or array, or dictionary, however, which is put away in a tenacious manner and controlled by a DBMS.

An efficient and compact structure of the index is used by the key-value store to have the option to rapidly and dependably find value using its key. For example, Redis is a key-value store used to tracklists, maps, heaps, and primitive types (which are simple data structures) in a constant database. Redis can uncover a very basic point of interaction to query and manipulate value types, just by supporting a predetermined number of value types, and when arranged, is prepared to do high throughput.

When to use a key-value database:

Here are a few situations in which you can use a key-value database:-

- User session attributes in an online app like finance or gaming, which is referred to as real-time random data access.
- Caching mechanism for repeatedly accessing data or key-based design.
- The application is developed on queries that are based on keys.

Features:

- One of the most un-complex kinds of NoSQL data models.
- For storing, getting, and removing data, key-value databases utilize simple functions.
- Querying language is not present in key-value databases.
- Built-in redundancy makes this database more reliable.

Advantages:

- It is very easy to use. Due to the simplicity of the database, data can accept any kind, or even different kinds when required.
- Its response time is fast due to its simplicity, given that the remaining environment near it is very much constructed and improved.
- Key-value store databases are scalable vertically as well as horizontally.
- Built-in redundancy makes this database more reliable.

Disadvantages:

- As querying language is not present in key-value databases, transportation of queries from one database to a different database cannot be done.
- The key-value store database is not refined. You cannot query the database without a key.

2.2.2 Document Data Model:

A Document Data Model is a lot different than other data models because it stores data in JSON, BSON, or XML documents. In this data model, we can move documents under one document and apart from this, any particular elements can be indexed to run queries faster. Often documents are stored and retrieved in such a way that it becomes close to the data objects which are used in many applications which means very less translations are required to use data in applications. JSON is a native language that is often used to store and query data too.

So in the document data model, each document has a key-value pair below is an example for the same.

```
{
  "Name" : "Yashodhra",
  "Address" : "Near Patel Nagar",
  "Email" : "yahoo123@yahoo.com",
  "Contact" : "12345"
}
```


Working of Document Data Model:

This is a data model which works as a semi-structured data model in which the records and data associated with them are stored in a single document which means this data model is not completely unstructured. The main thing is that data here is stored in a document.

- **Document Type Model:** As we all know data is stored in documents rather than tables or graphs, so it becomes easy to map things in many programming languages.
- **Flexible Schema:** Overall schema is very much flexible to support this statement one must know that not all documents in a collection need to have the same fields.
- **Distributed and Resilient:** Document data models are very much dispersed which is the reason behind horizontal scaling and distribution of data.
- **Manageable Query Language:** These data models are the ones in which query language allows the developers to perform CRUD (Create Read Update Destroy) operations on the data model.

Examples of Document Data Models :

- Amazon DocumentDB
- MongoDB
- Cosmos DB
- ArangoDB
- Couchbase Server
- CouchDB

Advantages:

- **Schema-less:** These are very good in retaining existing data at massive volumes because there are absolutely no restrictions in the format and the structure of data storage.
- **Faster creation of document and maintenance:** It is very simple to create a document and apart from this maintenance requires is almost nothing.
- **Open formats:** It has a very simple build process that uses XML, JSON, and its other forms.
- **Built-in versioning:** It has built-in versioning which means as the documents grow in size there might be a chance they can grow in complexity. Versioning decreases conflicts.

Disadvantages:

- **Weak Atomicity:** It lacks in supporting multi-document ACID transactions. A change in the document data model involving two collections will require us to run two separate queries i.e. one for each collection. This is where it breaks atomicity requirements.
- **Consistency Check Limitations:** One can search the collections and documents that are not connected to an author collection but doing this might create a problem in the performance of database performance.
- **Security:** Nowadays many web applications lack security which in turn results in the leakage of sensitive data. So it becomes a point of concern, one must pay attention to web app vulnerabilities.

Applications of Document Data Model :

- **Content Management:** These data models are very much used in creating various video streaming platforms, blogs, and similar services Because each is stored as a single document and the database here is much easier to maintain as the service evolves over time.
- **Book Database:** These are very much useful in making book databases because as we know this data model lets us nest.
- **Catalog:** When it comes to storing and reading catalog files these data models are very much used because it has a fast reading ability if incase Catalogs have thousands of attributes stored.
- **Analytics Platform:** These data models are very much used in the Analytics Platform.

2.2.3 Graph databases

A graph database is a type of NoSQL database that is designed to handle data with complex relationships and interconnections. In a graph database, data is stored as nodes and edges, where nodes represent entities and edges represent the relationships between those entities.

1. Graph databases are particularly well-suited for applications that require deep and complex queries, such as social networks, recommendation engines, and fraud detection systems. They can also be used for other types of applications, such as supply chain management, network and infrastructure management, and bioinformatics.
2. One of the main advantages of graph databases is their ability to handle and represent relationships between entities. This is because the relationships between entities are as important as the entities themselves, and often cannot be easily represented in a traditional relational database.
3. Another advantage of graph databases is their flexibility. Graph databases can handle data with changing structures and can be adapted to new use cases without requiring significant changes to the database schema. This makes them particularly useful for applications with rapidly changing data structures or complex data requirements.
4. However, graph databases may not be suitable for all applications. For example, they may not be the best choice for applications that require simple queries or that deal primarily with data that can be easily represented in a traditional relational database. Additionally, graph databases may require more specialized knowledge and expertise to use effectively.

The description of components are as follows:

- **Nodes:** represent the objects or instances. They are equivalent to a row in database. The node basically acts as a vertex in a graph. The nodes are grouped by applying a label to each member.
- **Relationships:** They are basically the edges in the graph. They have a specific direction, type and form patterns of the data. They basically establish relationship between nodes.
- **Properties:** They are the information associated with the nodes.

Some examples of Graph Databases software are Neo4j, Oracle NoSQL DB, Graph base etc. Out of which Neo4j is the most popular one.

In traditional databases, the relationships between data is not established. But in the case of Graph Database, the relationships between data are prioritized. Nowadays mostly interconnected data is used where one data is connected directly or indirectly. Since the concept of this database is based on graph theory, it is flexible and works very fast for associative data. Often data are interconnected to one another which also helps to establish further relationships. It works fast in the querying part as well because with the help of relationships we can quickly

find the desired nodes. join operations are not required in this database which reduces the cost. The relationships and properties are stored as first-class entities in Graph Database.

Graph databases allow organizations to connect the data with external sources as well. Since organizations require a huge amount of data, often it becomes cumbersome to store data in the form of tables. For instance, if the organization wants to find a particular data that is connected with another data in another table, so first join operation is performed between the tables, and then search for the data is done row by row. But Graph database solves this big problem. They store the relationships and properties along with the data. So if the organization needs to search for a particular data, then with the help of relationships and properties the nodes can be found without joining or without traversing row by row. Thus the searching of nodes is not dependent on the amount of data.

Types of Graph Databases:

- **Property Graphs:** These graphs are used for querying and analyzing data by modelling the relationships among the data. It comprises of vertices that has information about the particular subject and edges that denote the relationship. The vertices and edges have additional attributes called properties.
- **RDF Graphs:** It stands for Resource Description Framework. It focuses more on data integration. They are used to represent complex data with well defined semantics. It is represented by three elements: two vertices, an edge that reflect the subject, predicate and object of a sentence. Every vertex and edge is represented by URI(Uniform Resource Identifier).

When to Use Graph Database?

- Graph databases should be used for heavily interconnected data.
- It should be used when amount of data is larger and relationships are present.
- It can be used to represent the cohesive picture of the data.

How Graph and Graph Databases Work?

Graph databases provide graph models They allow users to perform traversal queries since data is connected. Graph algorithms are also applied to find patterns, paths and other relationships this enabling more analysis of the data. The algorithms help to explore the neighboring nodes, clustering of vertices analyze relationships and patterns. Countless joins are not required in this kind of database.

Example of Graph Database:

- Recommendation engines in E commerce use graph databases to provide customers with accurate recommendations, updates about new products thus increasing sales and satisfying the customer's desires.
- Social media companies use graph databases to find the "friends of friends" or products that the user's friends like and send suggestions accordingly to user.
- To detect fraud Graph databases play a major role. Users can create graph from the transactions between entities and store other important information. Once created, running a simple query will help to identify the fraud.

Advantages of Graph Database:

- Potential advantage of Graph Database is establishing the relationships with external sources as well
- No joins are required since relationships is already specified.
- Query is dependent on concrete relationships and not on the amount of data.
- It is flexible and agile.
- it is easy to manage the data in terms of graph.
- Efficient data modeling: Graph databases allow for efficient data modeling by representing data as nodes and edges. This allows for more flexible and scalable data modeling than traditional relational databases.
- Flexible relationships: Graph databases are designed to handle complex relationships and interconnections between data elements. This makes them well-suited for applications that require deep and complex queries, such as social networks, recommendation engines, and fraud detection systems.
- High performance: Graph databases are optimized for handling large and complex datasets, making them well-suited for applications that require high levels of performance and scalability.
- Scalability: Graph databases can be easily scaled horizontally, allowing additional servers to be added to the cluster to handle increased data volume or traffic.
- Easy to use: Graph databases are typically easier to use than traditional relational databases. They often have a simpler data model and query language, and can be easier to maintain and scale.

Disadvantages of Graph Database:

- Often for complex relationships speed becomes slower in searching.
- The query language is platform dependent.
- They are inappropriate for transactional data
- It has smaller user base.
- Limited use cases: Graph databases are not suitable for all applications. They may not be the best choice for applications that require simple queries or that deal primarily with data that can be easily represented in a traditional relational database.
- Specialized knowledge: Graph databases may require specialized knowledge and expertise to use effectively, including knowledge of graph theory and algorithms.
- Immature technology: The technology for graph databases is relatively new and still evolving, which means that it may not be as stable or well-supported as traditional relational databases.
- Integration with other tools: Graph databases may not be as well-integrated with other tools and systems as traditional relational databases, which can make it more difficult to use them in conjunction with other technologies.
- Overall, graph databases on NoSQL offer many advantages for applications that require complex and deep relationships between data elements. They are highly flexible, scalable, and performant, and can handle large and complex datasets. However, they may not be suitable for all applications, and may require specialized knowledge and expertise to use effectively.

Schemaless databases

Schemaless databases, also known as schema-free or schema-less databases, are a type of database management system (DBMS) that allows for flexible and dynamic data modeling without rigidly predefined schemas. In contrast to traditional relational databases that enforce a fixed schema, schemaless databases provide a more agile and adaptable approach for storing and querying data.

In schemaless databases, data is typically stored in a format that does not require a predefined schema to be specified before storing the data. Each document or record within the database can have its own structure, and the database system does not enforce a specific schema on the data. This means that different documents within the same collection or table can have varying structures and fields.

Some key characteristics of schemaless databases are as follows:

1. **Flexible Data Models:** Schemaless databases accommodate dynamic and evolving data structures. The absence of a predefined schema allows for easy addition, modification, or removal of fields within individual documents without affecting the entire database.
2. **No Fixed Relationships:** Schemaless databases do not enforce rigid relationships between entities or tables. Instead, relationships can be represented through embedded documents, references, or other techniques within the document structure itself.
3. **Agile Development:** Schemaless databases enable developers to iterate and modify the data model more rapidly since changes to the structure of individual documents can be made without requiring costly database migrations.
4. **Heterogeneous Data:** Schemaless databases can handle a wide variety of data types, including structured, semi-structured, and unstructured data. This makes them suitable for storing data in formats like JSON, BSON, XML, or other flexible formats.
5. **Scalability:** Schemaless databases are often designed for horizontal scalability. They can distribute data across multiple servers or clusters, allowing for high availability, performance, and scaling to handle large amounts of data.

Popular schemaless databases include MongoDB, Couchbase, and Apache Cassandra. These databases offer flexible data modeling capabilities, dynamic schema evolution, and scalable architectures suitable for various use cases, including content management systems, real-time analytics, and IoT applications.

How does a schemaless database work?

In schemaless databases, information is stored in JSON-style documents which can have varying sets of fields with different data types for each field. So, a collection could look like this:

```
{
  name : "Joe", age : 30, interests : 'football' }
{
  name : "Kate", age : 25
}
```

As you can see, the data itself normally has a fairly consistent structure. With the schemaless MongoDB database, there is some additional structure — the system namespace contains an explicit list of collections and indexes. Collections may be implicitly or explicitly created — indexes must be explicitly declared.

What are the benefits of using a schemaless database?

- **Greater flexibility over data types**

By operating without a schema, schemaless databases can store, retrieve, and query any data type — perfect for big data analytics and similar operations that are powered by unstructured data. Relational databases apply rigid schema rules to data, limiting what can be stored.

- **No pre-defined database schemas**

The lack of schema means that your NoSQL database can accept any data type — including those that you do not yet use. This future-proofs your database, allowing it to grow and change as your data-driven operations change and mature.

- **No data truncation**

A schemaless database makes almost no changes to your data; each item is saved in its own document with a partial schema, leaving the raw information untouched. This means that every detail is always available and nothing is stripped to match the current schema. This is particularly valuable if your analytics needs to change at some point in the future.

- **Suitable for real-time analytics functions**

With the ability to process unstructured data, applications built on NoSQL databases are better able to process real-time data, such as readings and measurements from IoT sensors. Schemaless databases are also ideal for use with machine learning and artificial intelligence operations, helping to accelerate automated actions in your business.

- **Enhanced scalability and flexibility**

With NoSQL, you can use whichever data model is best suited to the job. Graph databases allow you to view relationships between data points, or you can use traditional wide table views with an exceptionally large number of columns. You can query, report, and model information however you choose. And as your requirements grow, you can keep adding nodes to increase capacity and power.

2.6 Materialized views

Materialized views, also known as materialized or indexed views, are database objects that store the results of a query in a precomputed and persistent form. They are derived from one or more source tables or views and are used to improve query performance by providing faster access to frequently queried or complex data.

In a traditional database system, queries often involve joining multiple tables or performing complex calculations, which can be resource-intensive and time-consuming. Materialized views address this issue by precomputing and storing the results of such queries, allowing subsequent queries to retrieve the data directly from the materialized view instead of re-executing the original query.

1. **Data Storage:** Materialized views store the actual result set of a query, typically as a table-like structure in the database. The data in the materialized view is updated periodically to reflect changes in the underlying source tables.
2. **Query Performance:** By storing precomputed results, materialized views eliminate the need for executing complex queries repeatedly. This improves query performance by reducing the processing and computation time required for data retrieval.
3. **Data Aggregation and Joins:** Materialized views are commonly used for aggregating data or joining multiple tables to simplify and optimize complex queries. They can store the aggregated or joined results, allowing for faster access to the desired data.
4. **Maintenance and Refresh:** Materialized views need to be maintained and refreshed to reflect changes in the underlying data. Depending on the database system, materialized views can be refreshed on a schedule or triggered by specific events or updates to the source data.
5. **Query Rewrite:** Some database systems support automatic query rewrite, where the optimizer recognizes queries that can be satisfied using a materialized view and rewrites the query to use the materialized view instead. This further improves performance by transparently utilizing the materialized view.

Materialized views are particularly useful in scenarios where data is relatively static or updates are infrequent, and the benefits of improved query performance outweigh the overhead of maintaining the materialized view. They are commonly employed in data warehousing, reporting, and decision support systems where complex queries and aggregations are frequently performed.

It's important to note that the use of materialized views introduces some trade-offs. Materialized views require storage space to store the precomputed data, and there is a trade-off between storage space and query performance. Additionally, maintaining and refreshing materialized views can introduce overhead, especially in systems with frequent updates to the source data.

The availability and functionality of materialized views may vary across different database systems, so it's essential to consult the specific documentation and features of the database you are working with.

2.7 Distribution models

Distribution models in database systems refer to strategies for distributing data across multiple nodes or servers in a distributed computing environment. These models determine how data is partitioned and replicated to ensure availability, fault tolerance, and efficient query processing. Here are some commonly used distribution models:


1. **Horizontal Partitioning (Sharding):**
2. **Replication:**

The choice of a distribution model depends on factors such as the nature of the data, access patterns, scalability requirements, fault tolerance goals, and performance considerations. It's crucial to analyze the characteristics of the application and workload to determine the most suitable distribution model for a given scenario.

MongoDB, and Apache Hadoop, provide mechanisms to implement these distribution models and manage distributed data effectively.

2.7.1 Sharding

- Sharding involves dividing a large database into smaller, more manageable parts called shards or partitions.
- Each shard contains a subset of the data and is stored on a separate node or server in the distributed system.
- The sharding process typically involves selecting a shard key or partitioning key, which determines how data is distributed across shards.
- The goal of sharding is to evenly distribute data to avoid bottlenecks and enable horizontal scalability.
- Sharding is commonly used in NoSQL databases to handle large-scale datasets and achieve better performance and scalability.



Sharding example

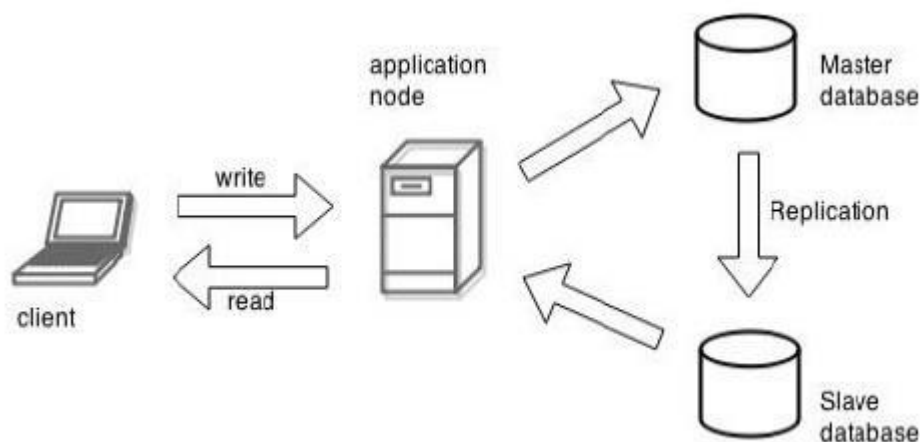
2.7.2 Replication

- Replication involves maintaining multiple copies (replicas) of data across different nodes in the distributed database cluster.
- Each replica is an exact copy of the data stored on a separate server.
- Replication enhances data availability, fault tolerance, and read performance by allowing data to be served from multiple replicas.
- Different replication models include master-slave replication and multi-master replication.
- In master-slave replication, one node (master) accepts write operations and asynchronously propagates changes to one or more replica nodes (slaves). Read queries can be distributed among replicas, reducing the load on the master.
- In multi-master replication, multiple nodes can accept write operations, and changes are replicated to other nodes. This approach allows for better write scaling and high availability.

Master-slave replication

Master-slave replication is a method of data replication in distributed database systems where one node, called the master or primary node, serves as the authoritative source for data, and one or more nodes, known as slave or secondary nodes, replicate and maintain copies of the master's data.

In a master-slave replication setup, the master node handles write operations (inserts, updates, deletes) and propagates those changes to the slave nodes. The slave nodes synchronize with the master to receive and apply the changes, ensuring that they have an up-to-date copy of the data. Slave nodes are typically read-only, meaning they do not accept write operations directly.



This data replication in NoSQL creates a copy (master copy) of your database and maintains it as the key data source. Any updates that you may require are made to this master copy and later transferred to the slave copies. Moreover, to maintain fast performance, all read requests are managed by the slave copies as it will not be feasible to burden the master copy alone. In case a master copy fails, one of the slave copies is automatically assigned as the new master.

Here are some key features and considerations of master-slave replication:

1. **Replication Process:** The master node logs all modifications made to the data, often in the form of binary logs or transaction logs. The slave nodes periodically connect to the master, retrieve the log files, and apply the logged changes to replicate the data.
2. **Data Consistency:** Master-slave replication ensures data consistency by replicating the changes made on the master to the slave nodes. By applying the same set of modifications in the same order, the slave nodes maintain a consistent copy of the data.
3. **Fault Tolerance and High Availability:** Master-slave replication provides fault tolerance and high availability. If the master node fails, one of the slave nodes can be promoted as the new master, ensuring that the system continues to operate with minimal disruption. This promotes data redundancy and reduces the risk of data loss.
4. **Read Scalability:** Slave nodes in master-slave replication can handle read queries, offloading the read traffic from the master node and distributing the workload across multiple nodes. This allows for horizontal scalability in read-intensive scenarios.

5. **Replication Lag:** There may be a delay between changes made on the master node and their replication to the slave nodes. This delay, known as replication lag, depends on factors such as network latency, load on the system, and the frequency of synchronization between the master and slave nodes. Applications that rely on real-time or near-real-time data consistency should consider the replication lag when accessing the slave nodes.
6. **Data Integrity:** Master-slave replication provides data integrity within the replicated data set. However, it does not protect against logical errors or data corruption on the master, as those changes would also be replicated to the slave nodes.

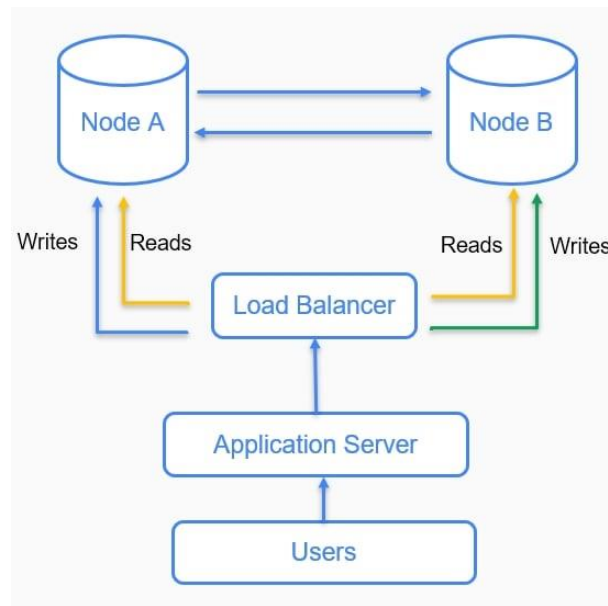
Master-slave replication is commonly used in distributed database systems to achieve data redundancy, fault tolerance, and improved read scalability. It is widely employed in various database technologies, including MySQL, PostgreSQL, and MongoDB, among others. However, it's important to note that master-slave replication is a single-master replication model, and more advanced replication models, such as multi-master or active-active replication, may be required for scenarios where multiple nodes can accept write operations simultaneously.

Peer-to-Peer Replication

In Peer-to-Peer Replication, there is no designated master node. Instead, all nodes in the database cluster are considered equal "peers." Each node can handle both read and write operations independently, and data is distributed across all nodes using techniques like sharding or consistent hashing to ensure horizontal scalability.

Characteristics of Peer-to-Peer Replication:

- No designated master node; all nodes are peers and can handle read and write operations.
- Data is distributed across all nodes in the cluster using techniques like sharding or consistent hashing.
- Provides better write scalability since write operations can be distributed across multiple nodes.
- Offers improved fault tolerance since there is no single point of failure like a master node.
- May require conflict resolution mechanisms for concurrent writes on different nodes.



2.8 Consistency

In NoSQL (Not Only SQL) databases, consistency refers to one of the four core principles known as the CAP theorem (Consistency, Availability, and Partition tolerance). The CAP theorem states that it is impossible for a distributed data store to simultaneously provide all three of these guarantees.

Specifically, in the context of NoSQL databases, consistency refers to the idea that all nodes in a distributed system see the same data at the same time. When a write operation is performed on one node, the data is immediately available and visible to all other nodes in the system. This ensures that clients accessing the database will always get the most recent data, providing a strong consistency model.

However, achieving strong consistency in a distributed system can come at the cost of increased latency and reduced availability. To ensure strong consistency, the system may need to wait for acknowledgments from all replicas before confirming a write operation, which can introduce delays. Additionally, if a network partition occurs (a portion of the nodes becomes unreachable), the system might have to sacrifice availability to maintain consistency.

Different NoSQL databases employ different consistency models based on their design goals and requirements. Some databases prioritize strong consistency, while others may choose to sacrifice it in favor of high availability or partition tolerance, depending on the use case.

Common consistency models in NoSQL databases include:

1. **Strong Consistency:** All nodes see the same data at the same time, ensuring the most up-to-date information. However, it might lead to slower response times and reduced availability during network partitions.
2. **Eventual Consistency:** The system guarantees that, given enough time and no further updates, all nodes will eventually converge to the same state. This model allows for improved availability and low-latency responses but might temporarily present inconsistent data during network partitions.

3. **Causal Consistency:** This model guarantees that causally related events (events that depend on each other) are observed in a consistent order across all nodes.
4. **Read-your-writes Consistency:** Ensures that if a client performs a write operation, it will always read back the most recent version of the data it wrote.

The choice of consistency model depends on the specific needs of the application and the trade-offs between consistency, availability, and partition tolerance that the system can afford. Different NoSQL databases offer varying degrees of consistency, and developers need to carefully consider these factors when selecting a database for their particular use case.

Relaxing consistency

Relaxing consistency, also known as eventual consistency, is a concept used in distributed systems and databases. It refers to a consistency model that allows for temporary inconsistencies between replicas or copies of data in a distributed environment.

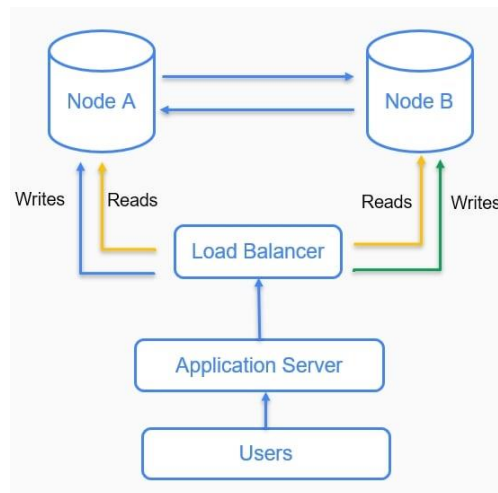
In an eventually consistent system:

1. **Data Updates:**
 - When a write operation occurs, the update is propagated asynchronously to different replicas in the system.
 - Each replica does not immediately reflect the update; instead, it continues to serve read requests based on its current state.
2. **Eventual Convergence:**
 - Over time, as the updates propagate and the system reconciles differences, all replicas will eventually converge to a consistent state.
 - This means that, given enough time and the absence of new updates, all replicas will eventually have the same data.
3. **Read Inconsistency:**
 - During the period of inconsistency, different replicas may show different views of the data.
 - Read operations might return stale data because the most recent updates have not yet reached all replicas.

Eventual consistency allows distributed systems to achieve high availability and fault tolerance, as it can continue to serve read and write requests even in the presence of network partitions or node failures. However, it comes at the cost of temporary inconsistency, and applications must be designed to handle this characteristic.

To work effectively with eventual consistency, applications often implement conflict resolution mechanisms to handle situations where concurrent updates result in conflicting data. By using version stamps, timestamps, or other techniques, conflicts can be resolved intelligently to maintain data integrity.

Eventual consistency is suitable for certain types of applications where immediate consistency is not strictly required, and eventual convergence to a consistent state is acceptable. It is commonly used in distributed databases, content delivery networks (CDNs), collaborative systems, and other scenarios where high availability and partition tolerance are critical.



2.9 The CAP theorem

The CAP theorem originally introduced as the CAP principle, can be used to explain some of the competing requirements in a distributed system with replication. It is a tool used to make system designers aware of the trade-offs while designing networked shared-data systems.

The three letters in CAP refer to three desirable properties of distributed systems with replicated data: **consistency** (among replicated copies), **availability** (of the system for read and write operations) and **partition tolerance** (in the face of the nodes in the system being partitioned by a network fault).

The CAP theorem states that it is not possible to guarantee all three of the desirable properties – consistency, availability, and partition tolerance at the same time in a distributed system with data replication.

The theorem states that networked shared-data systems can only strongly support two of the following three properties:

- **Consistency** – Consistency means that the nodes will have the same copies of a replicated data item visible for various transactions. A guarantee that every node in a distributed cluster returns the same, most recent and a successful write. Consistency refers to every client having the same view of the data. There are various types of consistency models. Consistency in CAP refers to sequential consistency, a very strong form of consistency.
- **Availability** – Availability means that each read or write request for a data item will either be processed successfully or will receive a message that the operation cannot be completed. Every non-failing node returns a response for all the read and write requests in a reasonable amount of time. The key word here is “every”. In simple terms, every node (on either side of a network partition) must be able to respond in a

reasonable amount of time.

- **Partition Tolerance –**

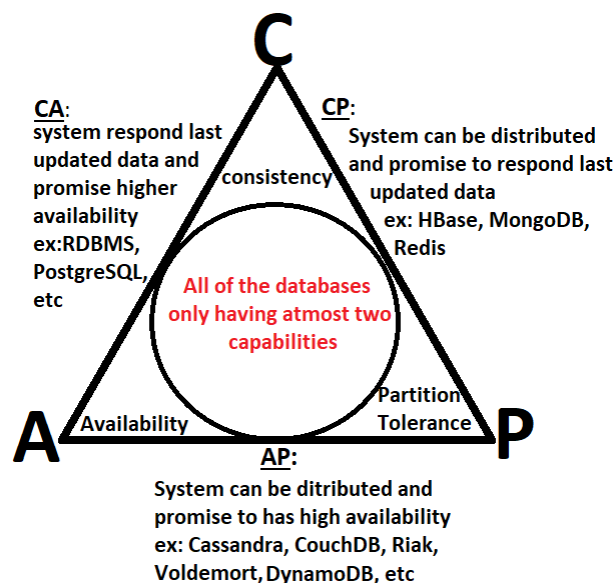
Partition tolerance means that the system can continue operating even if the network connecting the nodes has a fault that results in two or more partitions, where the nodes in each partition can only communicate among each other. That means, the system continues to function and upholds its consistency guarantees in spite of network partitions. Network partitions are a fact of life. Distributed systems guaranteeing partition tolerance can gracefully recover from partitions once the partition heals.

The use of the word consistency in CAP and its use in ACID do not refer to the same identical concept.

In CAP, the term consistency refers to the consistency of the values in different copies of the same data item in a replicated distributed system. In ACID, it refers to the fact that a transaction will not violate the integrity constraints specified on the database schema.

The CAP theorem states that distributed databases can have at most two of the three properties: consistency, availability, and partition tolerance. As a result, database systems prioritize only two properties at a time.

The following figure represents which database systems prioritize specific properties at a given time:



CAP theorem with databases examples

- **CA(Consistency and Availability)-**

The system prioritizes availability over consistency and can respond with possibly stale data.

Example databases: Cassandra, CouchDB, Riak, Voldemort.

- **AP(Availability and Partition Tolerance)-**

The system prioritizes availability over consistency and can respond with possibly stale data.

The system can be distributed across multiple nodes and is designed to operate reliably even in the face of network partitions.

Example databases: Amazon DynamoDB, Google Cloud Spanner.

- **CP(Consistency and Partition Tolerance)-**

The system prioritizes consistency over availability and responds with the latest updated data.

The system can be distributed across multiple nodes and is designed to operate reliably even in the face of network partitions.

Example databases: Apache HBase, MongoDB, Redis.

It's important to note that these database systems may have different configurations and settings that can change their behavior with respect to consistency, availability, and partition tolerance. Therefore, the exact behavior of a database system may depend on its configuration and usage.

for example, Neo4j, a graph database, the CAP theorem still applies. Neo4j prioritizes consistency and partition tolerance over availability, which means that in the event of a network partition or failure, Neo4j will sacrifice availability to maintain consistency.

2.10 Cassandra

Apache Cassandra is a highly scalable, distributed, and decentralized NoSQL database system designed to handle large amounts of data across multiple commodity servers or cloud instances. It provides high availability, fault tolerance, and linear scalability, making it suitable for use cases that require high write and read throughput with low latency.

Cassandra was initially developed at Facebook by two Indians Avinash Lakshman (one of the authors of Amazon's Dynamo) and Prashant Malik. It was developed to power the Facebook inbox search feature.

The following points specify the most important happenings in Cassandra history:

- Cassandra was developed at Facebook by Avinash Lakshman and Prashant Malik.
- It was developed for Facebook inbox search feature.
- It was open sourced by Facebook in July 2008.
- It was accepted by Apache Incubator in March 2009.
- Cassandra is a top level project of Apache since February 2010.
- The latest version of Apache Cassandra is 3.2.1.

Cassandra is commonly used in various applications and use cases, including real-time analytics, time-series data, IoT data management, recommendation engines, and content management systems. It is known for its ability to handle massive amounts of data with low latency and high scalability.

It's worth noting that Cassandra's distributed nature and tunable consistency come with trade-offs. Data replication and consistency require careful consideration, and the application needs to handle eventual consistency and conflict resolution. Additionally, data modeling in Cassandra may involve denormalization and designing tables based on query patterns to optimize performance.

Cassandra is open-source and managed by the Apache Software Foundation, with commercial distributions and cloud-managed services available as well.

Cassandra Architecture

Components of Cassandra

The main components of Cassandra are:


- **Node:** A Cassandra node is a place where data is stored.
- **Data center:** Data center is a collection of related nodes.
- **Cluster:** A cluster is a component which contains one or more data centers.
- **Commit log:** In Cassandra, the commit log is a crash-recovery mechanism. Every write operation is written to the commit log.
- **Mem-table:** A mem-table is a memory-resident data structure. After commit log, the data will be written to the mem-table. Sometimes, for a single-column family, there will be multiple mem-tables.
- **SSTable:** It is a disk file to which the data is flushed from the mem-table when its contents reach a threshold value.
- **Bloom filter:** These are nothing but quick, nondeterministic, algorithms for testing whether an element is a member of a set. It is a special kind of cache. Bloom filters are accessed after every query.

Some of the features of Cassandra architecture are as follows:

- Cassandra is designed such that it has no master or slave nodes.
- It has a ring-type architecture, that is, its nodes are logically distributed like a ring.
- Data is automatically distributed across all the nodes.
- Similar to HDFS, data is replicated across the nodes for redundancy.
- Data is kept in memory and lazily written to the disk.
- Hash values of the keys are used to distribute the data among nodes in the cluster.


A hash value is a number that maps any given key to a numeric value. For example, the string 'ABC' may be mapped to 101, and decimal number 25.34 may be mapped to 257. A hash value is generated using an algorithm so that the same value of the key always gives the same

hash value. In a ring architecture, each node is assigned a token value, as shown in the image below:



- Cassandra architecture supports multiple data centers.
- Data can be replicated across data centers.

You can keep three copies of data in one data center and the fourth copy in a remote data center for remote backup. Data reads prefer a local data center to a remote data center.



Cassandra Query Language

Cassandra Query Language (CQL) is used to access Cassandra through its nodes. CQL treats the database (Keyspace) as a container of tables. Programmers use `cqlsh`: a prompt to work with CQL or separate application language drivers.

The client can approach any of the nodes for their read-write operations. That node (coordinator) plays a proxy between the client and the nodes holding the data.


Cassandra Write Process

The Cassandra write process ensures fast writes. Steps in the Cassandra write process are:

1. Data is written to a commitlog on disk.

2. The data is sent to a responsible node based on the hash value.
3. Nodes write data to an in-memory table called memtable.
4. From the memtable, data is written to an sstable in memory. Sstable stands for Sorted String table. This has a consolidated data of all the updates to the table.
5. From the sstable, data is updated to the actual table.
6. If the responsible node is down, data will be written to another node identified as tempnode. The tempnode will hold the data temporarily till the responsible node comes alive.

The diagram below depicts the write process when data is written to table A.




Data is written to a commitlog on disk for persistence. It is also written to an in-memory memtable. Memtable data is written to sstable which is used to update the actual table.

Cassandra Read Process

The Cassandra read process ensures fast reads. Read happens across all nodes in parallel. If a node is down, data is read from the replica of the data. Priority for the replica is assigned on the basis of distance. Features of the Cassandra read process are:

- Data on the same node is given first preference and is considered data local.
- Data on the same rack is given second preference and is considered rack local.
- Data on the same data center is given third preference and is considered data center local.
- Data in a different data center is given the least preference.

Data in the memtable and sstable is checked first so that the data can be retrieved faster if it is already in memory. The diagram below represents a Cassandra cluster.



It has two data centers:

- data center 1
- data center 2


Data center 1 has two racks, while data center 2 has three racks. Fifteen nodes are distributed across this cluster with nodes 1 to 4 on rack 1, nodes 5 to 7 on rack 2, and so on.

Data Partitions

Cassandra performs transparent distribution of data by horizontally partitioning the data in the following manner:

- A hash value is calculated based on the primary key of the data.
- The hash value of the key is mapped to a node in the cluster
- The first copy of the data is stored on that node.
- The distribution is transparent as you can both calculate the hash value and determine where a particular row will be stored.

The following diagram depicts a four node cluster with token values of 0, 25, 50 and 75.



For a given key, a hash value is generated in the range of 1 to 100. Keys with hash values in the range 1 to 25 are stored on the first node, 26 to 50 are stored on the second node, 51 to 75 are stored on the third node, and 76 to 100 are stored on the fourth node. Please note that actual tokens and hash values in Cassandra are 127-bit positive integers.

Replication in Cassandra

Replication refers to the number of replicas that are maintained for each row. Replication provides redundancy of data for fault tolerance. A replication factor of 3 means that 3 copies of data are maintained in the system.

In this case, even if 2 machines are down, you can access your data from the third copy. The default replication factor is 1. A replication factor of 1 means that a single copy of the data is maintained, so if the node that has the data fails, you will lose the data.

Cassandra allows replication based on nodes, racks, and data centers, unlike HDFS that allows replication based on only nodes and racks. Replication across data centers guarantees data availability even when a data center is down.

2.11 Cassandra data model

The data model of Cassandra is significantly different from what we normally see in an RDBMS. This chapter provides an overview of how Cassandra stores its data.

Cluster

Cassandra database is distributed over several machines that operate together. The outermost container is known as the Cluster. For failure handling, every node contains a replica, and in case of a failure, the replica takes charge. Cassandra arranges the nodes in a cluster, in a ring format, and assigns data to them.

Keyspace

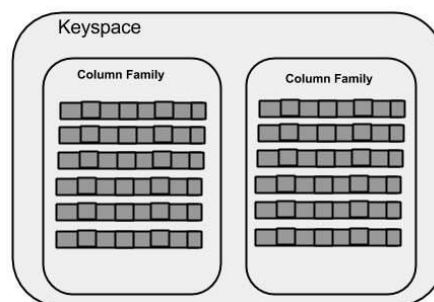
Keyspace is the outermost container for data in Cassandra. The basic attributes of a Keyspace in Cassandra are –

- **Replication factor** – It is the number of machines in the cluster that will receive copies of the same data.
- **Replica placement strategy** – It is nothing but the strategy to place replicas in the ring. We have strategies such as **simple strategy** (rack-aware strategy), **old network topology strategy** (rack-aware strategy), and **network topology strategy** (datacenter-shared strategy).
- **Column families** – Keyspace is a container for a list of one or more column families. A column family, in turn, is a container of a collection of rows. Each row contains ordered columns. Column families represent the structure of your data. Each keyspace has at least one and often many column families.

The syntax of creating a Keyspace is as follows –

```
CREATE KEYSPACE Keyspace name  
WITH replication = {'class': 'SimpleStrategy', 'replication_factor' : 3};
```

The following illustration shows a schematic view of a Keyspace.



Column Family

A column family is a container for an ordered collection of rows. Each row, in turn, is an ordered collection of columns. The following table lists the points that differentiate a column family from a table of relational databases.

Relational Table

A schema in a relational model is fixed. Once we define certain columns for a table, while inserting data, in every row all the columns must be filled at least with a null value.

Relational tables define only columns and the user fills in the table with values.

Cassandra column Family

In Cassandra, although the column families are defined, the columns are not. You can freely add any column to any column family at any time.


In Cassandra, a table contains columns, or can be defined as a super column family.

A Cassandra column family has the following attributes –

- **keys_cached** – It represents the number of locations to keep cached per SSTable.
- **rows_cached** – It represents the number of rows whose entire contents will be cached in memory.
- **preload_row_cache** – It specifies whether you want to pre-populate the row cache.

Note – Unlike relational tables where a column family’s schema is not fixed, Cassandra does not force individual rows to have all the columns.

The following figure shows an example of a Cassandra column family.



Column


A column is the basic data structure of Cassandra with three values, namely key or column name, value, and a time stamp. Given below is the structure of a column.

| Column | | |
|---------------|----------------|-----------------|
| name : byte[] | value : byte[] | clock : clock[] |

SuperColumn

A super column is a special column, therefore, it is also a key-value pair. But a super column stores a map of sub-columns.

Generally column families are stored on disk in individual files. Therefore, to optimize performance, it is important to keep columns that you are likely to query together in the same column family, and a super column can be helpful here. Given below is the structure of a super column.



Data Models of Cassandra and RDBMS

The following table lists down the points that differentiate the data model of Cassandra from that of an RDBMS.

| RDBMS | Cassandra |
|---|--|
| RDBMS deals with structured data. | Cassandra deals with unstructured data. |
| It has a fixed schema. | Cassandra has a flexible schema. |
| In RDBMS, a table is an array of arrays. (ROW x COLUMN) | In Cassandra, a table is a list of "nested key-value pairs". (ROW x COLUMN key x COLUMN value) |
| Database is the outermost container that contains data corresponding to an application. | Keyspace is the outermost container that contains data corresponding to an application. |
| Tables are the entities of a database. | Tables or column families are the entity of a keyspace. |
| Row is an individual record in RDBMS. | Row is a unit of replication in Cassandra. |
| Column represents the attributes of a relation. | Column is a unit of storage in Cassandra. |
| RDBMS supports the concepts of foreign keys, joins. | Relationships are represented using collections |

2.12 Cassandra example

Let's consider an example of a simple Cassandra data model for a hypothetical e-commerce application. We'll create a keyspace called "ecommerce" and define two tables: "users" and "products".

1. Keyspace: "ecommerce"

```
CREATE KEYSPACE ecommerce
WITH replication = {'class': 'SimpleStrategy', 'replication_factor': 3};
```

In this example, we use the SimpleStrategy replication strategy with a replication factor of 3, which means that each piece of data will be replicated to three nodes in the Cassandra cluster.

- Table: "users"

```
sql
CREATE TABLE ecommerce.users (
  user_id UUID PRIMARY KEY,
  name TEXT,
  email TEXT,
  address TEXT,
  phone_number TEXT
);
```

In the "users" table, we define the primary key as "user_id" of type UUID. The table also includes columns for the user's name, email, address, and phone number.

- Table: "products"

```
sql
CREATE TABLE ecommerce.products (
  product_id UUID PRIMARY KEY,
  name TEXT,
  price DECIMAL,
  description TEXT,
  category TEXT,
  in_stock BOOLEAN
);
```

The "products" table has a primary key defined on the "product_id" column of type UUID. Additionally, it includes columns for the product's name, price, description, category, and a boolean flag indicating whether the product is in stock.

This is a simple example of a Cassandra data model for an e-commerce application. You can expand and modify this model based on your specific requirements. Remember that in Cassandra, the data model should be designed based on the anticipated queries and access patterns to optimize for read performance and scalability.

Cassandra clients

Cassandra clients is used to connect, manage and develop your Cassandra database. The database client is used to manage your Cassandra database with actions like insert, delete, update table.

Cassandra provides support for multiple client drivers and APIs in various programming languages. These client drivers allow developers to interact with Cassandra and perform CRUD (Create, Read, Update, Delete) operations, execute queries, and manage the database. Here are some popular Cassandra client options:

1. **Java Driver:** The Java driver is the official client driver provided by the Apache Cassandra project. It offers high performance and comprehensive features for interacting with Cassandra from Java applications. It supports synchronous and asynchronous queries, automatic load balancing, connection pooling, and more.
2. **Python Driver:** The DataStax Python driver is a feature-rich and widely used driver for interacting with Cassandra from Python applications. It provides a Pythonic interface and supports features like automatic connection pooling, query building, and asynchronous execution. It is compatible with both Cassandra's native protocol and the Thrift API.
3. **C#/.NET Driver:** The DataStax C#/.NET driver is a robust and performant driver for accessing Cassandra from .NET applications. It supports both synchronous and asynchronous operations, connection pooling, automatic paging of large result sets, and advanced features like automatic query preparation and token-aware routing.
4. **Node.js Driver:** The DataStax Node.js driver allows developers to interact with Cassandra from Node.js applications. It offers a non-blocking, asynchronous API and supports features such as automatic query preparation, connection pooling, load balancing, and schema metadata retrieval.
5. **Ruby Driver:** The DataStax Ruby driver enables interaction with Cassandra from Ruby applications. It provides a simple and intuitive API, connection pooling, asynchronous operations, and automatic query preparation. It also includes support for Cassandra's lightweight transactions (LWTs) and batch statements.
6. **Go Driver:** The DataStax Go driver provides a native Go implementation for Cassandra connectivity. It supports concurrent query execution, connection pooling, automatic retry, and various consistency levels. It offers a convenient API for building and executing queries in a Go application.

These are just a few examples of the Cassandra client drivers available.