

Arnav fees upfront.?
data structure:
relational databases
were



ARNAV SAXENA

2K15/MC/014NoSQLNoSQL is a set of

database technologies that were developed in response to the demands presented in building modern applications (the early 2000s).

They are also called “ Not only SQL” to emphasize that they may support SQL- like query language too. It is an alternative to the traditional Relational Database Systems and can be used to store both structured as well as unstructured data. NoSQL vs RDBMS? Schema or No Schema: A NoSQL database lets you build an application without having to define the schema first, making it easier to update it as your data and requirements change unlike relational databases which turn rigid after forcing us to design the schema first.? Open Source: NoSQL databases are open source whereas relational databases typically are closed source with licensing fees baked into the use of their software. With NoSQL, you can get started on a project without any heavy investments in software fees upfront.? Data Structure: Relational databases were built in an era where data was fairly structured and clearly defined by their relationships.

NoSQL databases are designed to handle unstructured and semi-structured data (e. g., texts, social media posts, video, email) which makes up much of the data that exists today.? Scaling: NoSQL Databases employs the cheap and efficient Scale Out or Horizontal Scaling method as compared to the expensive Scale Up or Vertical Scaling used by Relational Databases. ? Development method: The long waterfall development cycle has been proved wrong time and again. Agile method is all the rage now where in teams keep iterating through the design quickly and pushing code every week or two, some even multiple times every day.

<https://assignbuster.com/arnav-fees-upfront-data-structure-relational-databases-were/>

Applications of NoSQL: Key Areas? Big Data Applications (eg. Hadoop, Spark)? Real-time Web Applications? Internet of Things? Mobile Applications? Content Management? Fraud Detection? Digital Communication

NoSQL Database Types Document databases pair each key with a complex data structure known as a document that encodes in itself information in formats like JSON. Common uses include managing content data monitoring Web and Mobile Applications. Eg. MongoDB, CouchDB.

Column Databases are optimized for queries over large datasets, and store columns of data together, instead of rows. Used for Internet Search and other large scale Web Applications. Eg. Apache Cassandra and HBase. Graph Database are used to store information about networks of data, such as social connections. Graph stores include Neo4J and Giraph. Key-value Databases: Every single item in the database is stored as an attribute name (or 'key'), together with its value, just like a hash table. It is used to maintain clickstream data, application logs etc.

Eg. Redis, Riak. Dynamicity: The core of NoSQL RDBMSes required us to define the schema first before entering any relevant data into our database. We had to conceive our tabular relations before hand. Many a times this schema would be rigid in nature and we would avoid updating it afterwards. This proved to be a hindrance in the way of Agile development where in the database often needs to be changed after regular iterations of development. Thus in a relational database in case we decided to update our schema at some point, we would have to create another database first and migrate our older data to the new schema.

For a large database, this migration process involves significant downtime and is pretty expensive. In case this happens frequently this downtime may keep on adding up and elongate the development cycle. Moreover, it's just not possible for relational databases to manage unstructured or semi-structured data that is being generated increasingly in today's 'Internet of Things' world. On the other hand, NoSQL databases allow data insertion without a predefined schema.

Any changes in the application can be made in real-time, without worrying about service interruptions - implying faster development and reliable code integration. Developers have typically had to add application-side code to enforce data quality controls, such as mandating the presence of specific fields, data types or permissible values. More sophisticated NoSQL databases allow validation rules to be applied within the database, allowing users to enforce governance across data, while maintaining the agility benefits of a dynamic schema. Relational Data Management in NoSQL: 1.

Collective querying Instead of retrieving all the data with one query, it is common to do several queries to get the desired data. NoSQL queries are often faster than traditional SQL queries so the cost of having to do additional queries may be acceptable. 2. Nesting In document databases we can put more data in a smaller number of collections. For instance, for a blog application, we may store comments within the blog post document so that with a single retrieval one gets all the comments. 3. Non-normalization Instead of only storing foreign keys, it is common to store actual foreign values along with the model's data. For example, each blog comment might

include the username in addition to a user id, thus providing easy access to the username without requiring another lookup.

When a username changes however, this will now need to be changed in many places in the database. Thus this approach works better when reads are much more common than writes. Limitations¹.

NoSQL being Open Source has a disadvantage of having different standards for different databases. 2. NoSQL is still in its beginning phase thereby making it difficult to find proper support or expert help. 3. NoSQL provides speed and scalability at the cost of ACID principles and inapt relational data handling. Conclusion NoSQL and RDBMS both have their uses and limitations.

Both are suited for different domains. Though NoSQL sounds to be a technology built to take over relational databases, but it must be noted that RDBMS enjoys its own exclusive advantages like maintaining referential integrity, ACID transactions and others. Thus in this Game of Data Management there's no single winner.